



Send documentation comments to mdsfeedback-doc@cisco.com.



Cisco MDS 9000 Family Configuration Guide, Release 2.x

Cisco MDS SAN-OS Release 2.0(1b) through Release 2.1(1a)
April 2005

Corporate Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

Text Part Number: OL-6973-02



THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

CCSP, CCVP, the Cisco Square Bridge logo, Follow Me Browsing, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Access Registrar, Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StrataView Plus, TeleRouter, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0502R)



CONTENTS

New and Changed Information xxxix

Preface xlv

| | |
|---|------|
| Audience | xlvi |
| Organization | xlvi |
| Document Conventions | xlvi |
| Related Documentation | xlix |
| Obtaining Documentation | i |
| Cisco.com | i |
| Documentation DVD | i |
| Ordering Documentation | i |
| Documentation Feedback | ii |
| Cisco Product Security Overview | ii |
| Reporting Security Problems in Cisco Products | ii |
| Obtaining Technical Assistance | iii |
| Cisco Technical Support Website | iii |
| Submitting a Service Request | iii |
| Definitions of Service Request Severity | iii |
| Obtaining Additional Publications and Information | liv |

CHAPTER 1

Product Overview 1-1

| | |
|---|-----|
| Hardware Overview | 1-1 |
| Cisco MDS 9100 Series Fixed Configuration Fabric Switches | 1-2 |
| Cisco MDS 9200 Series Fabric Switches | 1-2 |
| Cisco MDS 9500 Series Multilayer Directors | 1-3 |
| Software Features | 1-4 |
| Licensing | 1-4 |
| High Availability | 1-4 |
| Switch Reliability | 1-5 |
| Graceful Shut Down | 1-5 |
| Cisco Fabric Services | 1-5 |
| Virtual SANs | 1-6 |
| Dynamic VSANs | 1-6 |
| Intelligent Zoning | 1-7 |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|--|------|
| Enhanced Zoning | 1-7 |
| Device Alias Distribution | 1-7 |
| Inter-VSAN Routing | 1-7 |
| Trunking | 1-8 |
| PortChannels | 1-8 |
| IP Services | 1-8 |
| FICON | 1-9 |
| Fabric Binding | 1-9 |
| RLIR | 1-9 |
| IP Storage | 1-9 |
| Call Home | 1-10 |
| QoS and Congestion Control | 1-10 |
| SPAN and RSPAN | 1-11 |
| Switch Management Features | 1-11 |
| Redundant Supervisor Module Management | 1-11 |
| Fabric Management | 1-12 |
| Security Management | 1-12 |
| Port Tracking | 1-14 |
| SAN Extension Tuner | 1-14 |
| Command Scheduler | 1-14 |
| Intelligent Storage Services | 1-14 |
| Tools for Software Configuration | 1-14 |
| CLI | 1-15 |
| Cisco MDS 9000 Fabric Manager | 1-15 |

CHAPTER 2

Before You Begin 2-1

| | |
|--|------|
| About the Switch Prompt | 2-2 |
| Default Switch Roles | 2-3 |
| About the CLI Command Modes | 2-3 |
| CLI Command Hierarchy | 2-4 |
| EXEC Mode Options | 2-5 |
| Configuration Mode | 2-6 |
| Configuration Mode Commands and Submodes | 2-6 |
| CLI Command Navigation | 2-8 |
| Getting Help | 2-8 |
| Command Completion | 2-9 |
| File System Completion | 2-9 |
| The no and Default Forms of Commands | 2-9 |
| CLI Command Configuration Options | 2-10 |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|---|------|
| Displaying the Switch Configuration | 2-10 |
| Saving a Configuration | 2-13 |
| Clearing a Configuration | 2-13 |
| Displaying Users | 2-13 |
| Sending Messages to Users | 2-13 |
| Using the ping Command | 2-14 |
| Using the Extended ping Command | 2-14 |
| Using traceroute | 2-16 |
| Setting the Shell Timeout | 2-16 |
| Displaying VTY Sessions | 2-17 |
| Clearing VTY Sessions | 2-17 |
| Setting the Terminal Timeout | 2-17 |
| Setting the Terminal Type | 2-18 |
| Setting the Terminal Length | 2-18 |
| Setting the Terminal Width | 2-18 |
| Displaying Terminal Settings | 2-18 |
| Configuring the Switch Banner Message | 2-19 |
| About Flash Devices | 2-20 |
| Internal bootflash: | 2-20 |
| External CompactFlash (Slot0:) | 2-20 |
| Formatting Flash Devices and File Systems | 2-21 |
| Initializing Internal bootflash: | 2-21 |
| Formatting External CompactFlash | 2-21 |
| Using the File System | 2-22 |
| Setting the Current Directory | 2-22 |
| Displaying the Current Directory | 2-23 |
| Displaying File Checksums | 2-23 |
| Listing the Files in a Directory | 2-23 |
| Creating a Directory | 2-23 |
| Deleting an Existing Directory | 2-24 |
| Moving Files | 2-24 |
| Copying Files | 2-25 |
| Deleting Files | 2-25 |
| Displaying File Contents | 2-25 |
| Saving Command Output to a File | 2-26 |
| Compressing and Uncompressing Files | 2-26 |
| Displaying the Last Lines in a File | 2-27 |
| Executing Commands Specified in a Script | 2-27 |
| Setting the Delay Time | 2-28 |

Send documentation comments to mdsfeedback-doc@cisco.com.

CHAPTER 3
Obtaining and Installing Licenses 3-1

- Licensing Terminology 3-2
- Licensing Model 3-3
- Licensing High Availability 3-5
- Options to Install a License 3-6
- Obtaining a Factory-Installed License 3-6
- Performing a Manual Installation 3-6
- Obtaining the License Key File 3-7
- Installing the License Key File 3-8
- Backing Up License Files 3-9
- Identifying License Features in Use 3-9
- Uninstalling Licenses 3-9
- Updating Licenses 3-10
- License Expiry Alerts 3-11
- Grace Period Countdown 3-12
- License Transfers Between Switches 3-12
- Displaying License Information 3-12

CHAPTER 4
Initial Configuration 4-1

- Starting a Switch in the Cisco MDS 9000 Family 4-2
- Initial Setup Routine 4-2
 - Preparing to Configure the Switch 4-3
 - Default Login 4-3
 - Setup Options 4-4
 - Assigning Setup Information 4-5
 - Configuring Out-of-Band Management 4-5
 - Configuring In-Band Management 4-9
 - Using the setup Command 4-13
- Accessing the Switch 4-14
- Assigning a Switch Name 4-14
- Where Do You Go Next? 4-15
- Verifying the Module Status 4-15
- Configuring Date and Time 4-16
 - Configuring the Time Zone 4-16
 - Adjusting for Daylight Saving Time 4-16
 - NTP Configuration 4-17
 - NTP Configuration Guidelines 4-18

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|--|------|
| NTP Configuration Distribution | 4-19 |
| Committing NTP Configuration Changes | 4-20 |
| Discarding NTP Configuration Changes | 4-20 |
| Releasing Fabric Session Lock | 4-21 |
| Database Merge Guidelines | 4-21 |
| NTP Session Status Verification | 4-21 |
| Management Interface Configuration | 4-21 |
| Obtaining Remote Management Access | 4-22 |
| Using the force Option | 4-22 |
| Default Gateway Configuration | 4-23 |
| Configuring the Default Gateway | 4-23 |
| Telnet Server Connection | 4-24 |
| Disabling a Telnet Connection | 4-24 |
| Working with Configuration Files | 4-24 |
| Displaying Configuration Files | 4-25 |
| Downloading Configuration Files to the Switch | 4-25 |
| From a Remote Server | 4-26 |
| From an External CompactFlash Disk (slot0:) | 4-26 |
| Saving Configuration Files to an External Device | 4-26 |
| To a Remote Server | 4-27 |
| To an External CompactFlash Disk (slot0:) | 4-27 |
| Saving the Configuration | 4-27 |
| Copying Startup Configuration to the Fabric | 4-28 |
| Unlocking the Startup Configuration File | 4-28 |
| Copying Files | 4-28 |
| Backing Up the Current Configuration | 4-29 |
| Rolling Back to a Previous Configuration | 4-30 |
| Restoring the Configured Redundancy Mode | 4-30 |
| Downgrading from a Higher Release | 4-31 |
| Accessing Remote File Systems | 4-33 |
| Deleting Files | 4-33 |
| Configuring Console Port Settings | 4-34 |
| Verifying Console Port Settings | 4-34 |
| Configuring COM1 Port Settings | 4-35 |
| Verifying COM1 Port Settings | 4-35 |
| Configuring Modem Connections | 4-36 |
| Guidelines to Configure Modems | 4-36 |
| Enabling Modem Connections | 4-36 |
| Configuring the Initialization String | 4-37 |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|--|------|
| Configuring the Default Initialization String | 4-38 |
| Configuring a User-Specified Initialization String | 4-38 |
| Initializing a Modem in a Powered-On Switch | 4-39 |
| Verifying the Modem Connection Configuration | 4-39 |
| Configuring CDP | 4-40 |
| Clearing CDP Counters and Tables | 4-41 |
| Displaying CDP Information | 4-42 |

CHAPTER 5

Configuring High Availability 5-1

| | |
|--|-----|
| About High Availability | 5-1 |
| Switchover Mechanisms | 5-2 |
| HA Switchover Characteristics | 5-2 |
| Initiating a Switchover | 5-2 |
| Switchover Guidelines | 5-3 |
| Verifying Switchover Possibilities | 5-3 |
| Process Restartability | 5-4 |
| Synchronizing Supervisor Modules | 5-4 |
| Copying Boot Variable Images to the Standby Supervisor | 5-4 |
| Automatic Copying of Boot Variables | 5-4 |
| Verifying the Copied Boot Variables | 5-4 |
| Displaying HA Information | 5-5 |

CHAPTER 6

Software Images 6-1

| | |
|---|------|
| About Software Images | 6-1 |
| Dependent Factors for Software Installation | 6-2 |
| Essential Upgrade Prerequisites | 6-2 |
| Software Upgrade Methods | 6-4 |
| Determining Software Compatibility | 6-4 |
| Automated Upgrades | 6-5 |
| Benefits of Using the install all Command | 6-6 |
| Recognizing Failure Cases | 6-6 |
| Using the install all Command | 6-7 |
| Upgrading Services Modules | 6-9 |
| Sample install all Commands | 6-9 |
| Upgrade Status Verification | 6-17 |
| Manual Upgrade on a Dual Supervisor Switch | 6-18 |
| Preparing for a Manual Installation | 6-18 |
| Upgrading a Loader | 6-19 |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|--|------|
| Upgrading the BIOS | 6-21 |
| Quick Upgrade | 6-23 |
| Maintaining Supervisor Modules | 6-23 |
| Standby Supervisor Boot Variable Version | 6-23 |
| Standby Supervisor Bootflash Memory | 6-24 |
| Standby Supervisor Boot Alert | 6-24 |
| Replacing Modules | 6-24 |
| Corrupted Bootflash Recovery | 6-25 |
| Recovery Using BIOS Setup | 6-27 |
| Recovery from the loader> Prompt | 6-30 |
| Recovery from the switch(boot)# Prompt | 6-31 |
| Recovery for Switches with Dual Supervisor Modules | 6-32 |
| Recognizing Error States | 6-33 |
| Default Settings | 6-34 |

CHAPTER 7

Managing Modules 7-1

| | |
|---|------|
| About Modules | 7-1 |
| Supervisor Modules | 7-2 |
| Switching Modules | 7-3 |
| Services Modules | 7-3 |
| Verifying the Status of a Module | 7-3 |
| Checking the State of a Module | 7-4 |
| Connecting to a Module | 7-4 |
| Reloading Modules | 7-6 |
| Reloading the Switch | 7-6 |
| Power Cycling Modules | 7-6 |
| Reloading Switching Modules | 7-6 |
| Preserving Module Configuration | 7-7 |
| Purging Module Configuration | 7-8 |
| Powering Off Switching Modules | 7-8 |
| Identifying Module LEDs | 7-9 |
| EPLD Configuration | 7-13 |
| Upgrading EPLD Images | 7-13 |
| Displaying EPLD Versions | 7-14 |
| Installing the ASM and Specifying the ASM Image Boot Variable | 7-15 |
| Configuring the ASM-SFN Image Boot Variable for VSFN | 7-15 |
| Configuring the SSI Image Boot Variable for Fibre Channel Switching | 7-17 |
| Installing the SSM and Specifying the SSI Image Boot Variable | 7-18 |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|--|------|
| Configuring the ASM-SFN Image Boot Variable for VSFN | 7-19 |
| Configuring the SSI Image Boot Variable for Fibre Channel Switching and Intelligent Storage Services | 7-21 |
| Replacing Modules | 7-22 |
| Default Settings | 7-23 |

CHAPTER 8

Managing System Hardware 8-1

| | |
|---------------------------------------|------|
| Displaying Switch Hardware Inventory | 8-2 |
| Displaying the Switch Serial Number | 8-4 |
| Displaying Power Usage Information | 8-5 |
| Power Supply Configuration Modes | 8-6 |
| Power Supply Configuration Guidelines | 8-6 |
| About Module Temperature | 8-9 |
| Displaying Module Temperature | 8-10 |
| About Fan Modules | 8-10 |
| About Clock Modules | 8-11 |
| Displaying Environment Information | 8-11 |
| Default Settings | 8-12 |

CHAPTER 9

Using the CFS Infrastructure 9-1

| | |
|--|-----|
| About CFS | 9-2 |
| Cisco SAN-OS Features Using CFS | 9-2 |
| CFS Features | 9-3 |
| CFS Protocol | 9-3 |
| CFS Distribution Scopes | 9-4 |
| CFS Distribution Modes | 9-4 |
| Uncoordinated Distribution | 9-4 |
| Coordinated Distribution | 9-4 |
| Disabling CFS Distribution on a Switch | 9-5 |
| CFS Application Requirements | 9-5 |
| Enabling CFS for an Application | 9-5 |
| Locking the Fabric | 9-6 |
| Committing Changes | 9-6 |
| Discarding Changes | 9-6 |
| Saving the Configuration | 9-6 |
| Clearing a Locked Session | 9-7 |
| CFS Merge Support | 9-7 |

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying CFS Configuration Information 9-7

Default Settings 9-11

CHAPTER 10

Configuring and Managing VSANs 10-1

VSAN Advantages 10-1

How VSANs Work 10-2

VSANs Versus Zones 10-4

Default and Isolated VSANs 10-5

Default VSAN 10-5

Isolated VSAN 10-5

Displaying Isolated VSAN Membership 10-5

VSAN Attributes 10-6

Operational State of a VSAN 10-6

VSAN Membership 10-6

Creating and Configuring VSANs Statically 10-7

Assigning Static Port VSAN Membership 10-7

Deleting Static VSANs 10-8

Displaying Static VSAN Configurations 10-9

Default Settings 10-10

CHAPTER 11

Creating Dynamic VSANs 11-1

About DPVM 11-2

DPVM Requirements 11-2

Enabling DPVM 11-2

About DPVM Databases 11-3

Configuring Config and Pending Databases 11-3

Activating Config Databases 11-3

About Autolearned Entries 11-4

Enabling Autolearning 11-4

Clearing Learned Entries 11-5

Configuring DPVM Database Distribution 11-5

Disabling DPVM Database Distribution 11-5

Locking the Fabric 11-6

Committing Changes 11-6

Discarding Changes 11-6

Clearing a Locked Session 11-7

Database Merge Guidelines 11-7

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|--------------------------------|-------|
| Copying DPVM Databases | 11-7 |
| Comparing Database Differences | 11-8 |
| Displaying DPVM Configurations | 11-8 |
| Sample DPVM Configuration | 11-10 |
| Default Settings | 11-12 |

CHAPTER 12

Configuring Interfaces 12-1

| | |
|-------------------------------------|-------|
| Fibre Channel Interfaces | 12-2 |
| About Interface Modes | 12-3 |
| E Port | 12-3 |
| F Port | 12-4 |
| FL Port | 12-4 |
| TL Port | 12-4 |
| TE Port | 12-4 |
| SD Port | 12-5 |
| ST Port | 12-5 |
| Fx Port | 12-5 |
| B Port | 12-5 |
| Auto Mode | 12-5 |
| About Interface States | 12-6 |
| Administrative States | 12-6 |
| Operational States | 12-6 |
| Reason Codes | 12-6 |
| Configuring Fibre Channel Interface | 12-9 |
| Graceful Shut Down | 12-9 |
| Interface Modes | 12-10 |
| TL Port ALPA Caches | 12-10 |
| Displaying the ALPA Cache Contents | 12-11 |
| Clearing the ALPA Cache | 12-11 |
| Administrative Speeds | 12-11 |
| Interface Descriptions | 12-12 |
| Buffer-to-Buffer Credits | 12-12 |
| Performance Buffers | 12-13 |
| Extended BB_credits | 12-14 |
| Frame Encapsulation | 12-16 |
| Receive Data Field Size | 12-16 |
| Beacon Mode | 12-16 |
| Identifying the Beacon LEDs | 12-17 |
| About Speed LEDs | 12-17 |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|--------------------------------------|-------|
| Switch Port Attribute Default Values | 12-17 |
| SFP Transmitter Types | 12-18 |
| Configuring Management Interfaces | 12-19 |
| Configuring VSAN Interfaces | 12-20 |
| Configuring CIM | 12-20 |
| Added Security on a CIM Server | 12-20 |
| Displaying Interface Information | 12-21 |
| Displaying TL Port Information | 12-30 |
| TL Port Translation Guidelines | 12-31 |
| Default Settings | 12-33 |

CHAPTER 13

Configuring Trunking 13-1

| | |
|---|------|
| About Trunking | 13-1 |
| About the Trunking Protocol | 13-2 |
| Enabling or Disabling the Trunking Protocol | 13-2 |
| Configuring Trunk Mode | 13-2 |
| Configuring the Trunk Mode | 13-3 |
| Trunk-Allowed VSAN Configuration | 13-3 |
| Configuring an Allowed-Active List of VSANs | 13-5 |
| Trunking Configuration Guidelines | 13-6 |
| Displaying Trunking Information | 13-7 |
| Default Settings | 13-8 |

CHAPTER 14

Configuring PortChannels 14-1

| | |
|---|-------|
| PortChannel Functionality | 14-2 |
| PortChannel Examples | 14-2 |
| 32-Port Switching Module Configuration Guidelines | 14-3 |
| About PortChanneling and Trunking | 14-4 |
| About Load Balancing | 14-4 |
| PortChannel Creation | 14-6 |
| PortChannel Modes | 14-7 |
| Deleting PortChannels | 14-8 |
| Interface Addition to a PortChannel | 14-8 |
| Forcing an Interface Addition | 14-9 |
| Compatibility Check | 14-10 |
| Suspended and Isolated States | 14-10 |
| Deleting Interfaces from a PortChannel | 14-11 |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|--|-------|
| PortChannel Configuration Guidelines | 14-11 |
| Error Detection | 14-11 |
| Valid Configurations | 14-12 |
| Invalid Configuration Examples | 14-13 |
| PortChannel Protocol | 14-13 |
| About PortChannel Protocols | 14-14 |
| Channel Group Creation | 14-14 |
| Autocreation Functionality | 14-15 |
| Enabling and Configuring Autocreation | 14-16 |
| Converting to Manually-Configured Channel Groups | 14-16 |
| PortChannel Configuration Verification | 14-17 |
| Default Settings | 14-20 |

CHAPTER 15

| | |
|--|-------------|
| Configuring and Managing Zones | 15-1 |
| Zoning Features | 15-2 |
| Zoning Example | 15-3 |
| Zone Implementation | 15-4 |
| Zone Configuration | 15-4 |
| Configuring a Zone | 15-5 |
| Alias Configuration | 15-6 |
| Zone Set Creation | 15-7 |
| Active and Full Zone Set Considerations | 15-8 |
| Activating a Zone Set | 15-10 |
| Zone Enforcement | 15-10 |
| The Default Zone | 15-11 |
| Zone Set Distribution | 15-11 |
| Enabling Full Zone Set Distribution | 15-12 |
| One-Time Distribution | 15-12 |
| Recovering from Link Isolation | 15-13 |
| Importing and Exporting Zone Sets | 15-14 |
| Zone Set Duplication | 15-14 |
| Copying Zone Sets | 15-14 |
| Zone Database Information | 15-15 |
| Clearing the Zone Server Database | 15-15 |
| Zone-Based Traffic Priority | 15-15 |
| Configuring Default Zone QoS Priority Attributes | 15-16 |
| Configuring Broadcast Zoning | 15-17 |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|---|-------|
| About LUN Zoning | 15-17 |
| Configuring a LUN-Based Zone | 15-19 |
| Assigning LUNs to Storage Subsystems | 15-19 |
| About Read-Only Zones | 15-19 |
| Read-Only Zone Configuration Guidelines | 15-19 |
| Configuring Read-Only Zones | 15-20 |
| Renaming Zones, Zone Sets, fcaliases, and Zone Attribute Groups | 15-21 |
| Cloning Zones, Zone Sets, fcaliases, and Zone Attribute Groups | 15-21 |
| Displaying Zone Information | 15-21 |
| About Enhanced Zoning | 15-27 |
| Advantages of Enhanced Zoning | 15-28 |
| Changing from Basic Zoning to Enhanced Zoning | 15-28 |
| Changing from Enhanced Zoning to Basic Zoning | 15-29 |
| Enabling Enhanced Zoning | 15-29 |
| Modifying the Zone Database | 15-30 |
| Creating Attribute Groups | 15-30 |
| Merging the Database | 15-31 |
| The Merge Process | 15-31 |
| Default Zone Policies | 15-32 |
| Broadcasting a Zone | 15-32 |
| Displaying Enhanced Zone Information | 15-33 |
| Default Settings | 15-36 |

CHAPTER 16

| | |
|--|-------------|
| Distributing Device Alias Services | 16-1 |
| About Device Aliases | 16-1 |
| Device Alias Features | 16-2 |
| Device Alias Requirements | 16-2 |
| Zone Aliases Versus Device Aliases | 16-2 |
| Modifying the Device Alias Database | 16-3 |
| Locking The Fabric | 16-3 |
| Committing Changes | 16-4 |
| Discarding Changes | 16-4 |
| Fabric Lock Override | 16-4 |
| Device Alias Distribution | 16-5 |
| Legacy Zone Alias Configuration Conversion | 16-5 |
| Database Merge Guidelines | 16-5 |
| Device Alias Statistics Cleanup | 16-6 |
| Device Alias Configuration Verification | 16-6 |

Send documentation comments to mdsfeedback-doc@cisco.com.

Default Settings 16-10

CHAPTER 17

Configuring Inter-VSAN Routing 17-1

About IVR 17-2

IVR Features 17-3

IVR Terminology 17-3

IVR Guidelines 17-4

Domain ID Guidelines 17-4

Transit VSAN Guidelines 17-5

Border Switch Guidelines 17-5

IVR Configuration 17-5

Unique Domain ID Configuration Options 17-6

Enabling IVR 17-6

IVR Configuration Distribution 17-6

Database Implementation 17-7

Enabling Configuration Distribution 17-7

Locking the Fabric 17-7

Committing the Changes 17-7

Discarding the Changes 17-8

Clearing a Locked Session 17-8

About IVR NAT 17-8

Enabling IVR NAT 17-9

About IVR Topologies 17-10

Configuring IVR Topologies 17-10

Manually Configuring the IVR Topology 17-10

Configuring an IVR Topology Database 17-11

Activating a Manually Configured IVR Topology 17-12

Configuring IVR Topology Automatic Mode 17-13

Migrating from IVR Topology Automatic Mode to Manual Mode 17-13

Clearing the Configured IVR Topology Database 17-14

Non-Unique VSAN IDs Using AFIDs 17-14

Configuring the AFID Database 17-14

Adding IVR Virtual Domain 17-15

About IVZs and IVZSs 17-16

IVZs Versus Zones 17-16

Automatic IVZ Creation 17-16

Configuring IVZs and IVZSs 17-18

Creating and Activating IVZs and IVZSs 17-18

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|--|-------|
| Configuring LUNs in IVR Zoning | 17-19 |
| Configuring the QoS Attribute | 17-19 |
| Using the force Option | 17-20 |
| Clearing the IVZ Database | 17-20 |
| Displaying IVZ Configurations | 17-21 |
| About IVR Service Groups | 17-21 |
| Configuring IVR Service Groups | 17-22 |
| IVR Interoperability | 17-22 |
| Configuring IVR Using Read-Only Zoning | 17-22 |
| Database Merge Guidelines | 17-23 |
| Configuring IVR Logging Levels | 17-25 |
| Displaying IVR Information | 17-25 |
| Sample Configuration | 17-28 |
| Default Settings | 17-31 |

CHAPTER 18

Managing FLOGI, Name Server, FDMI, and RSCN Databases 18-1

| | |
|---|------|
| Displaying FLOGI Details | 18-1 |
| About the Name Server Proxy Feature | 18-2 |
| Registering Name Server Proxies | 18-3 |
| Rejecting Duplicate pWWN | 18-3 |
| Displaying Name Server Database Entries | 18-3 |
| Displaying FDMI | 18-5 |
| About RSCN Information | 18-7 |
| Displaying RSCN Information | 18-7 |
| About the multi-pid Option | 18-8 |
| Clearing RSCN Statistics | 18-9 |

CHAPTER 19

Configuring Switch Security 19-1

| | |
|----------------------------------|------|
| Switch Management Security | 19-2 |
| CLI Security Options | 19-2 |
| SNMP Security Options | 19-2 |
| Switch AAA Functionalities | 19-2 |
| Authentication | 19-2 |
| Authorization | 19-3 |
| Accounting | 19-3 |
| Remote AAA Services | 19-4 |
| Remote Authentication Guidelines | 19-4 |
| Server Groups | 19-4 |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|--|-------|
| AAA Service Configuration Options | 19-4 |
| Error-Enabled Status | 19-5 |
| Configuring RADIUS | 19-5 |
| Setting the RADIUS Server Address | 19-6 |
| Setting the Global Preshared Key | 19-6 |
| Setting the RADIUS Server Timeout Interval | 19-7 |
| Setting Iterations of the RADIUS Server | 19-7 |
| Defining Vendor-Specific Attributes | 19-7 |
| VSA Format | 19-8 |
| Specifying SNMPv3 on AAA Servers | 19-8 |
| Displaying RADIUS Server Details | 19-9 |
| Configuring TACACS+ | 19-10 |
| About TACACS+ | 19-10 |
| Enabling TACACS+ | 19-10 |
| Setting the TACACS+ Server Address | 19-10 |
| Setting the Global Secret Key | 19-11 |
| Setting the Timeout Value | 19-12 |
| Defining Custom Attributes for Roles | 19-12 |
| Supported TACACS+ Servers | 19-13 |
| Displaying TACACS+ Server Details | 19-13 |
| Configuring Server Groups | 19-14 |
| Distributing AAA Server Configuration | 19-15 |
| Enabling the RADIUS Server Distribution | 19-15 |
| Starting a Distribution Session on a Switch | 19-16 |
| Displaying the Session Status | 19-16 |
| Displaying the Configuration to Be Distributed | 19-17 |
| Committing the Distribution | 19-17 |
| Discarding the Distribution Session | 19-17 |
| Clearing Sessions | 19-17 |
| Merge Guidelines for RADIUS and TACACS+ Configurations | 19-18 |
| Local AAA Services | 19-19 |
| Disabling AAA Authentication | 19-19 |
| Displaying AAA Authentication | 19-19 |
| Authentication and Authorization Process | 19-20 |
| Role-Based Authorization | 19-21 |
| Configuring Roles and Profiles | 19-22 |
| Configuring Rules and Features for Each Role | 19-23 |
| Modifying Profiles | 19-23 |
| Configuring the VSAN Policy | 19-24 |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|--|-------|
| Modifying the VSAN Policy | 19-24 |
| Distributing Role-Based Configurations | 19-25 |
| Database Implementation | 19-25 |
| Locking The Fabric | 19-25 |
| Committing the Changes | 19-25 |
| Discarding the Changes | 19-25 |
| Enabling Distribution | 19-26 |
| Clearing Sessions | 19-26 |
| Database Merge Guidelines | 19-26 |
| Displaying Role-Based Information | 19-26 |
| Displaying Role-Based When Distribution is Enabled | 19-27 |
| Configuring User Accounts | 19-29 |
| Creating or Updating Users | 19-29 |
| Logging out Users | 19-31 |
| Displaying User Account Information | 19-31 |
| SNMP Security | 19-32 |
| Configuring Accounting Services | 19-32 |
| Displaying Accounting Configuration | 19-32 |
| Clearing Accounting Logs | 19-34 |
| Configuring SSH Services | 19-34 |
| Enabling SSH Service | 19-34 |
| Specifying the SSH Key | 19-34 |
| Generating the SSH Server Key Pair | 19-35 |
| Overwriting a Generated Key Pair | 19-35 |
| Clearing SSH Hosts | 19-36 |
| Displaying SSH Protocol Status | 19-37 |
| Recovering Administrator Password | 19-37 |
| Configuring Cisco ACS Servers | 19-38 |
| Default Settings | 19-42 |

CHAPTER 20

| | |
|---|-------------|
| Configuring Fabric Security | 20-1 |
| About Fabric Authentication | 20-2 |
| About DHCHAP | 20-3 |
| DHCHAP Compatibility with Existing Cisco MDS Features | 20-3 |
| Configuring DHCHAP Authentication | 20-3 |
| DHCHAP Configuration | 20-4 |
| DHCHAP Authentication Modes | 20-4 |
| DHCHAP Hash Algorithm Configuration | 20-5 |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|--|-------|
| DHCHAP Group Configuration | 20-6 |
| DHCHAP Password Configuration | 20-6 |
| Configuring the DHCHAP Password for the Local Switch | 20-7 |
| Password Configuration for Other Devices | 20-7 |
| Locally Configuring the Device Name | 20-8 |
| DHCHAP Timeout Value | 20-8 |
| Configuring the Timeout Value | 20-8 |
| Displaying Protocol Security Information | 20-9 |
| DHCHAP AAA Authentication | 20-10 |
| Sample Configuration | 20-10 |
| Default Settings | 20-12 |

CHAPTER 21

| | |
|--|-------------|
| Configuring Port Security | 21-1 |
| Port Security Features | 21-2 |
| Port Security Enforcement | 21-2 |
| Port Security Initiation | 21-2 |
| Port Security Manual Configuration | 21-3 |
| WWN Identification | 21-3 |
| Authorized Port Pair Addition | 21-3 |
| Port Security Activation | 21-4 |
| Database Activation Rejection | 21-5 |
| Forceful Port Security Activation | 21-5 |
| Database Reactivation | 21-6 |
| About AutoLearning | 21-7 |
| Configuring auto-learn | 21-7 |
| Disabling Autolearning | 21-7 |
| Auto-Learning Device Authorization | 21-8 |
| Authorization Scenario | 21-8 |
| Port Security Configuration Distribution | 21-9 |
| Enabling Distribution | 21-9 |
| Locking The Fabric | 21-10 |
| Committing the Changes | 21-10 |
| Discarding the Changes | 21-10 |
| Activation and Autolearning Configuration Distribution | 21-10 |
| Database Merge Guidelines | 21-12 |
| Database Interaction | 21-12 |
| Database Scenarios | 21-12 |
| Port Security Database Copy | 21-13 |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|---|-------|
| Port Security Database Deletion | 21-14 |
| Port Security Database Cleanup | 21-14 |
| Displaying Port Security Configurations | 21-15 |
| Default Settings | 21-18 |

CHAPTER 22

| | |
|--|-------------|
| Configuring SNMP | 22-1 |
| SNMP Security | 22-2 |
| SNMP Version 1 and Version 2c | 22-2 |
| SNMP Version 3 | 22-2 |
| SNMPv3 CLI User Management and AAA Integration | 22-3 |
| CLI and SNMP User Synchronization | 22-3 |
| Restricting Switch Access | 22-3 |
| Group-Based SNMP Access | 22-4 |
| Configuring Common Roles | 22-4 |
| Mapping of CLI operations to SNMP | 22-5 |
| Creating and Modifying Users | 22-6 |
| Configuring SNMP Users from the CLI | 22-6 |
| Enforcing SNMPv3 Message Encryption | 22-7 |
| Assigning SNMPv3 Users to Multiple Roles | 22-8 |
| AES Encryption-Based Privacy | 22-8 |
| Adding or Deleting Communities | 22-9 |
| Assigning SNMP Switch Contact Information | 22-9 |
| Configuring SNMP Notifications (Traps and Informs) | 22-9 |
| Configuring SNMPv1 and SNMPv2c Notifications | 22-10 |
| Configuring SNMPv3 Notifications | 22-10 |
| Enabling SNMP Notifications | 22-11 |
| Configuring the Notification Target User | 22-12 |
| Displaying SNMP Security Information | 22-13 |
| Default Settings | 22-15 |

CHAPTER 23

| | |
|--------------------------|-------------|
| Configuring RMON | 23-1 |
| About RMON | 23-1 |
| Configuring RMON | 23-1 |
| RMON Alarm Configuration | 23-2 |
| RMON Event Configuration | 23-3 |
| RMON Verification | 23-3 |
| Default Settings | 23-3 |

Send documentation comments to mdsfeedback-doc@cisco.com.

CHAPTER 24
Configuring Fibre Channel Routing Services and Protocols 24-1

| | |
|---|-------|
| FSPF Features | 24-2 |
| FSPF Examples | 24-2 |
| Fault Tolerant Fabric | 24-2 |
| Redundant Links | 24-3 |
| Fail-Over Scenarios for PortChannels and FSPF Links | 24-3 |
| FSPF Global Configuration | 24-4 |
| Global FSPF Configuration | 24-4 |
| FSPF Configuration Deletion | 24-5 |
| FSPF Routing Protocol Usage | 24-5 |
| Link State Record Defaults | 24-5 |
| FSPF Interface Configuration | 24-6 |
| FSPF Link Cost | 24-6 |
| Hello Time Intervals | 24-6 |
| Dead Time Intervals | 24-6 |
| Disabling FSPF for Specific Interfaces | 24-7 |
| Retransmitting Intervals | 24-7 |
| Configuring Fibre Channel Routes | 24-8 |
| Clearing FSPF Counters | 24-9 |
| Broadcast and Multicast Routing | 24-10 |
| In-Order Delivery | 24-11 |
| Reordering Network Frames | 24-11 |
| Reordering PortChannel Frames | 24-11 |
| Enabling In-Order Delivery | 24-12 |
| Enabling IOD Globally | 24-12 |
| Enabling IOD for a VSAN | 24-13 |
| Displaying the IOD Status | 24-13 |
| Configuring the Drop Latency Time | 24-14 |
| Displaying Latency Information | 24-14 |
| Flow Statistics Configuration | 24-15 |
| Configuring Flow Statistics | 24-15 |
| Counting Flow Statistics | 24-15 |
| Clearing FIB Statistics | 24-16 |
| Displaying Flow Statistics | 24-16 |
| Displaying Routing and Forwarding Information | 24-17 |
| Displaying Global FSPF Information | 24-19 |
| Displaying the FSPF Database | 24-20 |
| Displaying FSPF Interfaces | 24-21 |

Send documentation comments to mdsfeedback-doc@cisco.com.

Default Settings 24-21

CHAPTER 25

Configuring Intelligent Storage Services 25-1

- About SCSI Flow Services 25-2
 - SCSI Flow Manager 25-3
 - SCSI Flow Configuration Client 25-3
 - SCSI Flow Data Path Support 25-3
- Configuring SCSI Flow Services 25-3
 - Enabling SCSI Flow Services 25-4
 - Enabling SCSI Flow Configuration Distribution 25-4
 - Configuring SCSI Flow Identifiers 25-5
- About Fibre Channel Write Acceleration 25-5
- Enabling Fibre Channel Write Acceleration 25-5
- About SCSI Flow Statistics 25-6
- Enabling SCSI Flow Statistics 25-7
 - Clearing SCSI Flow Statistics 25-7
- Displaying SCSI Flow Services Information 25-7
- About SANTap 25-10
- Enabling SANTap 25-14
- Displaying SANTap Information 25-15
- About NASB 25-17
- Enabling NASB 25-19
- Displaying NASB Information 25-20
- Default Settings 25-21

CHAPTER 26

Configuring IP Services 26-1

- Traffic Management Services 26-2
- Management Interface Configuration 26-2
- Default Gateway Configuration 26-3
- Default Network Configuration 26-3
- IP Access Control Lists 26-5
 - IP-ACL Configuration Guidelines 26-5
 - Filter Contents 26-5
 - Protocol Information 26-5
 - Address Information 26-6
 - Port Information 26-6
 - ICMP Information 26-7

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|--|-------|
| TOS Information | 26-7 |
| IP-ACL -Creation | 26-8 |
| Adding filters to an Existing IP-ACL | 26-9 |
| Removing Entries from an Existing IP-ACL | 26-9 |
| Reading the IP-ACL Log Dump | 26-9 |
| Applying an IP-ACL to an Interface | 26-10 |
| IP-ACL Configuration Verification | 26-11 |
| IP-ACL Counter Cleanup | 26-11 |
| IPFC Configuration | 26-12 |
| Configuring an IP Address in a VSAN | 26-12 |
| Enabling IP Routing | 26-12 |
| Configuring IP Static Routes | 26-13 |
| Displaying and Clearing ARPs | 26-13 |
| Displaying IP Interface Information | 26-14 |
| Overlay VSAN Configuration | 26-15 |
| Multiple VSAN Configuration | 26-17 |
| The Virtual Router Redundancy Protocol | 26-19 |
| VRRP Features | 26-19 |
| VRRP Functionality | 26-19 |
| Virtual Router Addition and Deletion | 26-20 |
| Virtual Router Initiation | 26-21 |
| Virtual Router IP Address Addition | 26-21 |
| Priority for the Virtual Router | 26-21 |
| Time Interval for Advertisement Packets | 26-22 |
| Virtual Router Authentication | 26-23 |
| Priority Based on Interface State | 26-23 |
| Displaying VRRP Information | 26-24 |
| Clearing VRRP Statistics | 26-25 |
| DNS Server Configuration | 26-25 |
| Displaying DNS Host Information | 26-26 |
| Default Settings | 26-26 |

CHAPTER 27

Configuring FICON 27-1

| | |
|--------------------------------|------|
| About FICON | 27-2 |
| FICON Requirements | 27-2 |
| MDS-Specific FICON Advantages | 27-3 |
| Fabric Optimization with VSANs | 27-3 |
| FCIP Support | 27-4 |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|--|-------|
| PortChannel Support | 27-4 |
| VSANs for FICON and FCP Intermixing | 27-4 |
| Cisco MDS-Supported FICON Features | 27-5 |
| FICON Port Numbering | 27-7 |
| Port Addresses | 27-8 |
| Implemented and Unimplemented Port Addresses | 27-8 |
| Installed and Uninstalled Ports | 27-9 |
| FICON Port Numbering Guidelines | 27-9 |
| FCIP and PortChannel Port Numbers | 27-9 |
| FC ID Allocation | 27-10 |
| FICON Cascading | 27-10 |
| FICON VSAN Prerequisites | 27-11 |
| Enabling FICON | 27-11 |
| Effects of Enabling FICON | 27-11 |
| Setting Up a Basic FICON Configuration | 27-12 |
| Manually Enabling FICON | 27-15 |
| The code-page Option | 27-16 |
| FC ID Last Byte | 27-16 |
| FICON Host Control | 27-17 |
| Host Moves the Switch Offline | 27-17 |
| Host Changes FICON Port Parameters | 27-17 |
| Host Controls the Time Stamp | 27-17 |
| Time Stamp Cleanup | 27-18 |
| FICON SNMP Control | 27-18 |
| Running Configuration Automatic Save | 27-19 |
| Binding Port Numbers to PortChannels | 27-20 |
| Binding Port Numbers to FCIP Interfaces | 27-20 |
| Configuring FICON Ports | 27-21 |
| Port Blocking | 27-21 |
| Port Prohibiting | 27-21 |
| Port Address Name Assignment | 27-22 |
| FICON Configuration Files | 27-23 |
| Accessing FICON Configuration Files | 27-23 |
| Applying the FICON Configuration Files | 27-24 |
| Editing FICON Configuration Files | 27-24 |
| Copying FICON Configuration Files | 27-25 |
| Port Swapping | 27-25 |
| Port Swapping Guidelines | 27-26 |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|---|-------|
| Moving a FICON VSAN to an Offline State | 27-27 |
| Clearing FICON Device Allegiance | 27-27 |
| CUP In-Band Management | 27-27 |
| Placing CUPs in a Zone | 27-28 |
| Displaying FICON Information | 27-28 |
| Receiving FICON Alerts | 27-28 |
| Displaying FICON Port Address Information | 27-29 |
| Displaying IPL File Information | 27-30 |
| Displaying the Configured FICON State | 27-32 |
| Displaying a Ports Administrative State | 27-33 |
| Displaying Control Unit Information | 27-33 |
| Displaying Buffer Information | 27-34 |
| Displaying FICON Information in the Running Configuration | 27-35 |
| Displaying FICON Information in the Startup Configuration | 27-36 |
| Displaying FICON-Related Log Information | 27-36 |
| Fabric Binding Configuration | 27-37 |
| Port Security Versus Fabric Binding | 27-37 |
| Fabric Binding Enforcement | 27-38 |
| Fabric Binding Initiation | 27-38 |
| Switch WWN List Configuration | 27-39 |
| Fabric Binding Activation | 27-39 |
| Forcing Fabric Binding Activation | 27-40 |
| Saving Fabric Binding Configurations | 27-40 |
| Clearing the Fabric Binding Statistics | 27-41 |
| Deleting the Fabric Binding Database | 27-41 |
| Verifying Fabric Binding Configurations | 27-41 |
| Displaying RLIR Information | 27-44 |
| Clearing RLIR Information | 27-48 |
| Default Settings | 27-48 |

CHAPTER 28

| | |
|--|-------------|
| Configuring IP Storage | 28-1 |
| Services Modules | 28-2 |
| Module Status Verification | 28-3 |
| IPS Module Upgrade | 28-4 |
| MPS-14/2 Module Upgrade | 28-4 |
| Supported Hardware | 28-4 |
| Configuring Gigabit Ethernet Interfaces | 28-4 |
| Configuring a Basic Gigabit Ethernet Interface | 28-5 |
| Configuring Interface Descriptions | 28-5 |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|--|-------|
| Configuring Beacon Mode | 28-5 |
| Configuring Auto-Negotiation | 28-6 |
| Configuring the MTU Frame Size | 28-6 |
| Configuring Promiscuous Mode | 28-6 |
| About VLANs for Gigabit Ethernet | 28-7 |
| Configuring the VLAN Subinterface | 28-7 |
| Interface Subnet Requirements | 28-8 |
| Configuring Static IP Routing | 28-8 |
| Displaying the IP Route Table | 28-9 |
| Verifying Gigabit Ethernet Connectivity | 28-9 |
| Gigabit Ethernet IP-ACL Guidelines | 28-10 |
| Applying IP-ACLs on Gigabit Ethernet Interfaces | 28-10 |
| Displaying ARP Caches | 28-11 |
| Clearing ARP Caches | 28-11 |
| Displaying Statistics | 28-12 |
| Displaying Gigabit Ethernet Interface Statistics | 28-12 |
| Displaying Ethernet MAC Statistics | 28-12 |
| Displaying DMA-Bridge Statistics | 28-13 |
| Displaying TCP/IP Statistics | 28-13 |
| Configuring Gigabit Ethernet High Availability | 28-15 |
| VRRP for iSCSI and FCIP Services | 28-16 |
| Configuring VRRP for Gigabit Ethernet Interfaces | 28-16 |
| About Ethernet PortChannel Aggregation | 28-17 |
| Configuring Ethernet PortChannels | 28-18 |
| Configuring CDP | 28-19 |
| Configuring FCIP | 28-19 |
| FCIP and VE Ports | 28-20 |
| FCIP Links | 28-21 |
| FCIP Profiles | 28-21 |
| FCIP Interfaces | 28-22 |
| Enabling FCIP | 28-22 |
| Basic FCIP Configuration | 28-22 |
| Creating FCIP Profiles | 28-23 |
| Creating FCIP Links | 28-24 |
| Advanced FCIP Profile Configuration | 28-25 |
| Configuring TCP Listener Ports | 28-25 |
| Configuring TCP Parameters | 28-25 |
| Advanced FCIP Interface Configuration | 28-30 |
| Configuring Peers | 28-30 |
| Active Connections | 28-32 |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|---|-------|
| Number of TCP Connections | 28-32 |
| Time Stamp Control | 28-32 |
| B Port Interoperability Mode | 28-34 |
| Quality of Service | 28-36 |
| Configuring E Ports | 28-36 |
| Advanced FCIP Features | 28-37 |
| FCIP Write Acceleration | 28-37 |
| FCIP Tape Acceleration | 28-39 |
| FCIP Compression | 28-41 |
| Displaying FCIP Information | 28-42 |
| FCIP High Availability | 28-48 |
| Fibre Channel PortChannels | 28-48 |
| FSPF | 28-49 |
| VRRP | 28-49 |
| Ethernet PortChannels | 28-50 |
| Ethernet PortChannels and Fibre Channel PortChannels | 28-50 |
| Configuring iSCSI | 28-51 |
| Enabling iSCSI | 28-53 |
| Creating iSCSI Interfaces | 28-53 |
| Presenting Fibre Channel Targets as iSCSI Targets | 28-54 |
| Dynamic Mapping | 28-54 |
| Static Mapping | 28-56 |
| iSCSI Virtual Target Configuration Examples | 28-57 |
| Presenting iSCSI Hosts as Virtual Fibre Channel Hosts | 28-59 |
| Initiator Identification | 28-59 |
| Initiator Presentation Modes | 28-60 |
| VSAN Membership for iSCSI | 28-65 |
| Example of VSAN membership for iSCSI devices | 28-66 |
| Advanced VSAN membership for iSCSI hosts | 28-67 |
| iSCSI Access Control | 28-67 |
| Fibre Channel Zoning Based Access Control | 28-67 |
| iSCSI ACL Based Access Control | 28-69 |
| Enforcing Access Control | 28-70 |
| iSCSI Session Authentication | 28-70 |
| Authentication Mechanism | 28-71 |
| Local Authentication | 28-71 |
| Restricting iSCSI Initiator Authentication | 28-72 |
| Mutual CHAP Authentication | 28-72 |
| iSCSI Interface Advanced Features | 28-74 |
| iSCSI Listener Port | 28-74 |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|---|--------|
| TCP Tuning parameters | 28-74 |
| QoS | 28-74 |
| iSCSI Routing Modes | 28-75 |
| Displaying iSCSI Information | 28-77 |
| Displaying iSCSI Interfaces | 28-77 |
| Displaying iSCSI Statistics | 28-77 |
| Displaying Proxy Initiator Information | 28-79 |
| Displaying Global iSCSI Information | 28-80 |
| Displaying iSCSI Sessions | 28-81 |
| Displaying iSCSI Initiators | 28-82 |
| Displaying iSCSI Virtual Targets | 28-86 |
| Displaying iSCSI User Information | 28-86 |
| iSCSI High Availability | 28-86 |
| Transparent Target Failover | 28-86 |
| Multiple IPS Ports Connected to the Same IP Network | 28-90 |
| VRRP-Based High Availability | 28-92 |
| Ethernet PortChannel-Based High Availability | 28-93 |
| iSCSI Authentication Setup Guidelines and Scenarios | 28-93 |
| No Authentication | 28-93 |
| CHAP with Local Password Database | 28-94 |
| CHAP with External RADIUS Server | 28-94 |
| iSCSI Transparent Mode Initiator | 28-95 |
| Target Storage Device Requiring LUN Mapping | 28-101 |
| Configuring iSCSI Storage Name Services | 28-106 |
| iSNS Client Functionality | 28-107 |
| Creating an iSNS Profile | 28-107 |
| Verifying iSNS Client Configuration | 28-108 |
| iSNS Server Functionality | 28-110 |
| Example Scenario | 28-110 |
| Enabling the iSNS Server | 28-112 |
| iSCSI Configuration Distribution | 28-112 |
| ESI Retry Count Configuration | 28-112 |
| iSNS Client Registration and Deregistration | 28-113 |
| Target Discovery | 28-113 |
| Verifying the iSNS Server Configuration | 28-114 |
| IPS Module Core Dumps | 28-120 |
| Default Settings | 28-121 |

Send documentation comments to mdsfeedback-doc@cisco.com.

CHAPTER 29

Configuring IPsec Network Security 29-1

| | |
|---|-------|
| About IPsec | 29-2 |
| About IKE | 29-3 |
| IPsec Prerequisites | 29-3 |
| IPsec Compatibility | 29-4 |
| IPsec and IKE Terminology | 29-5 |
| Supported IPsec Transforms and Algorithms | 29-6 |
| Supported IKE Transforms and Algorithms | 29-6 |
| Initializing IKE | 29-7 |
| Configuring the IKE Domain | 29-7 |
| About IKE Tunnels | 29-8 |
| IKE Policy Negotiation | 29-8 |
| Optional Configurations | 29-10 |
| Clearing IKE Tunnels or Domains | 29-11 |
| Refreshing SAs | 29-12 |
| Configuring IPsec | 29-12 |
| Crypto ACLs | 29-12 |
| Crypto ACL Guidelines | 29-13 |
| Mirror Image Crypto ACLs | 29-14 |
| The any Keyword in Crypto ACLs | 29-16 |
| Transform Sets in IPsec | 29-16 |
| Crypto Map Entries | 29-18 |
| SA Establishment Between Peers | 29-18 |
| Crypto Map Configuration Guidelines | 29-19 |
| SA Lifetime Negotiation | 29-19 |
| The auto-peer Option | 29-20 |
| Perfect Forward Secrecy | 29-22 |
| Crypto Map Set Interface Application | 29-22 |
| IPsec Maintenance | 29-23 |
| Global Lifetime Values | 29-23 |
| Displaying IKE Configurations | 29-24 |
| Displaying IPsec Configurations | 29-25 |
| Sample FCIP Configuration | 29-30 |
| Sample iSCSI Configuration | 29-34 |
| Default Settings | 29-36 |

Send documentation comments to mdsfeedback-doc@cisco.com.

CHAPTER 30

| | |
|---|-------------|
| Configuring Call Home | 30-1 |
| Call Home Features | 30-2 |
| Cisco AutoNotify | 30-2 |
| Call Home Configuration Process | 30-3 |
| Contact Information | 30-3 |
| Destination Profiles | 30-4 |
| Alert Groups | 30-6 |
| Call Home Message Levels | 30-8 |
| Syslog-based Alerts | 30-9 |
| RMON-based Alerts | 30-9 |
| E-Mail Options | 30-10 |
| Configuring General E-Mail Options | 30-10 |
| Configuring SMTP Server and Ports | 30-10 |
| Periodic Inventory Notification | 30-11 |
| Duplicate Message Throttle | 30-11 |
| Call Home Enable Function | 30-12 |
| Call Home Configuration Distribution | 30-12 |
| Fabric Lock Override | 30-13 |
| Database Merge Guidelines | 30-14 |
| Call Home Communications Test | 30-14 |
| Displaying Call Home Information | 30-14 |
| Sample Syslog Alert Notification in Full-txt Format | 30-16 |
| Sample Syslog Alert Notification in XML Format | 30-17 |
| Sample RMON Notification in XML Format | 30-17 |
| Default Settings | 30-18 |
| Event Triggers | 30-19 |
| Call Home Message Levels | 30-20 |
| Message Contents | 30-21 |

CHAPTER 31

| | |
|--------------------------------------|-------------|
| Configuring Domain Parameters | 31-1 |
| About fcdomain Phases | 31-2 |
| Domain Restart | 31-3 |
| Domain Configuration | 31-4 |
| Switch Priority | 31-6 |
| Allowed Domain ID Lists | 31-6 |
| Merged Stable Fabrics | 31-7 |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|--|-------|
| Contiguous Domain Assignments | 31-7 |
| fcdomain Initiation | 31-8 |
| Fabric Name | 31-8 |
| Incoming RCFs | 31-9 |
| Persistent FC IDs | 31-9 |
| Persistent FC IDs Manual Configuration | 31-10 |
| Unique Area FC IDs for Some HBAs | 31-11 |
| Persistent FC ID Selective Purging | 31-13 |
| Displaying fcdomain Information | 31-13 |
| Default Settings | 31-17 |

CHAPTER 32

Configuring Traffic Management 32-1

| | |
|--|-------|
| FCC | 32-2 |
| FCC Process | 32-2 |
| Enabling FCC | 32-3 |
| Assigning FCC Priority | 32-3 |
| Displaying FCC | 32-3 |
| QoS | 32-3 |
| Control Traffic | 32-4 |
| Disabling Control Traffic | 32-4 |
| Displaying Control Traffic Information | 32-4 |
| Data Traffic | 32-4 |
| VSAN Versus Zone-Based QoS | 32-6 |
| Configuring Data Traffic | 32-6 |
| QoS Initiation for Data Traffic | 32-6 |
| Class Map Creation | 32-7 |
| Service Policy Definition | 32-8 |
| Service Policy Enforcement | 32-9 |
| DWRR Traffic Scheduler | 32-9 |
| Displaying Data Traffic Information | 32-10 |
| Example Configuration | 32-12 |
| Ingress Port Rate Limiting | 32-13 |
| Default Settings | 32-14 |

CHAPTER 33

Tracking and Redirecting Traffic 33-1

| | |
|---------------------------|------|
| About Port Tracking | 33-2 |
| Port Tracking Terminology | 33-2 |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|--------------------------------------|------|
| Port Tracking Guidelines | 33-3 |
| Port Tracking Features | 33-3 |
| Enabling Port Tracking | 33-3 |
| Configuring Linked Ports | 33-3 |
| Operational Binding | 33-4 |
| Tracking Multiple Ports | 33-4 |
| Monitoring Ports in a VSAN | 33-5 |
| Forceful Shutdown | 33-5 |
| Displaying Port Tracking Information | 33-6 |
| Default Settings | 33-8 |

CHAPTER 34

Configuring the SAN Extension Tuner 34-1

| | |
|-----------------------------------|------|
| About SET | 34-2 |
| License Prerequisites | 34-2 |
| Tuner Guidelines | 34-3 |
| Tuner Initialization | 34-3 |
| Tuner Configuration | 34-3 |
| nWWN Configuration | 34-4 |
| Virtual N Port Configuration | 34-5 |
| SCSI Read/Write Assignment | 34-5 |
| Data Pattern | 34-7 |
| Tuning Configuration Verification | 34-8 |
| Default Settings | 34-9 |

CHAPTER 35

Scheduling Maintenance Jobs 35-1

| | |
|----------------------------------|------|
| About the Command Scheduler | 35-1 |
| Scheduler Terminology | 35-1 |
| Scheduling Guidelines | 35-2 |
| Scheduler Configuration | 35-2 |
| Command Scheduler Initialization | 35-3 |
| Job Definition | 35-3 |
| Job Deletion | 35-4 |
| Schedule Definition | 35-5 |
| Periodic Schedule Definition | 35-5 |
| One-Time Schedule Definition | 35-6 |
| Schedule Deletion | 35-6 |
| Job Disassociation | 35-7 |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|--------------------------------------|------|
| Schedule Time Deletion | 35-7 |
| Execution Log | 35-7 |
| Clearing the Log File Contents | 35-7 |
| Scheduler Configuration Verification | 35-8 |
| Default Settings | 35-9 |

CHAPTER 36

| | |
|---|-------------|
| Configuring System Message Logging | 36-1 |
| About System Message Logging | 36-1 |
| System Message Logging Configuration | 36-3 |
| Message Logging Initiation | 36-4 |
| Console Severity Level | 36-4 |
| Module Logging | 36-4 |
| Facility Severity Level | 36-5 |
| Log Files | 36-5 |
| System Message Logging Servers | 36-6 |
| Outgoing System Message Logging Server Facilities | 36-6 |
| System Message Logging Configuration Distribution | 36-7 |
| Fabric Lock Override | 36-8 |
| Database Merge Guidelines | 36-8 |
| Displaying System Message Logging Information | 36-9 |
| Default Settings | 36-13 |

CHAPTER 37

| | |
|---------------------------------|-------------|
| Discovering SCSI Targets | 37-1 |
| About SCSI LUN Discovery | 37-1 |
| Starting SCSI LUN Discovery | 37-2 |
| Initiating Customized Discovery | 37-2 |
| Displaying SCSI LUN Information | 37-3 |

CHAPTER 38

| | |
|--|-------------|
| Monitoring Network Traffic Using SPAN | 38-1 |
| About SPAN | 38-2 |
| SPAN Sources | 38-3 |
| IPS Source Ports | 38-3 |
| CSM Source Ports | 38-4 |
| Allowed Source Interface Types | 38-4 |
| VSAN as a Source | 38-4 |
| Guidelines to Configure VSANs as a Source | 38-4 |
| SPAN Sessions | 38-5 |
| Specifying Filters | 38-6 |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|--|-------|
| Guidelines to Specifying Filters | 38-6 |
| SD Port Characteristics | 38-6 |
| Guidelines to Configure SPAN | 38-6 |
| Configuring SPAN | 38-7 |
| Encapsulating Frames | 38-8 |
| SPAN Conversion Behavior | 38-9 |
| Monitoring Traffic Using Fibre Channel Analyzers | 38-10 |
| Without SPAN | 38-10 |
| With SPAN | 38-11 |
| Configuring Analyzers Using SPAN | 38-11 |
| Single SD Port to Monitor Traffic | 38-12 |
| Displaying SPAN Information | 38-13 |
| Remote SPAN | 38-15 |
| Advantages to Using RSPAN | 38-15 |
| FC and RSPAN Tunnels | 38-16 |
| Guidelines to Configure RSPAN | 38-16 |
| ST Port Characteristics | 38-17 |
| Configuring RSPAN | 38-17 |
| RSPAN Configuration Example | 38-17 |
| Configuration in the Source Switch | 38-17 |
| Configuration in All Intermediate Switches | 38-20 |
| Configuration in the Destination Switch | 38-21 |
| Explicit Paths | 38-23 |
| Monitoring RSPAN Traffic | 38-25 |
| Sample Scenarios | 38-25 |
| Single Source with One RSPAN Tunnel | 38-26 |
| Single Source with Multiple RSPAN Tunnels | 38-26 |
| Multiple Sources with Multiple RSPAN Tunnels | 38-27 |
| Displaying RSPAN Information | 38-27 |
| Default SPAN and RSPAN Settings | 38-29 |
| Default RSPAN Settings | 38-30 |

CHAPTER 39

Advanced Features and Concepts 39-1

| | |
|--------------------------------------|------|
| Fibre Channel Time Out Values | 39-2 |
| Timer Configuration Across All VSANs | 39-2 |
| Timer Configuration Per-VSAN | 39-2 |
| fctimer Distribution | 39-3 |
| Committing fctimer Changes | 39-3 |
| Discarding fctimer Changes | 39-4 |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|--|-------|
| Fabric Lock Override | 39-4 |
| Database Merge Guidelines | 39-4 |
| Displaying Configured FC Timer Values | 39-5 |
| The fctrace Feature | 39-5 |
| The fcping Feature | 39-7 |
| Verifying Switch Connectivity | 39-8 |
| Configuring a Fabric Analyzer | 39-8 |
| About the Cisco Fabric Analyzer | 39-9 |
| Local Text-Based Capture | 39-9 |
| Remote Capture Daemon | 39-10 |
| GUI-Based Client | 39-10 |
| Configuring the Cisco Fabric Analyzer | 39-10 |
| Capturing Frames Locally | 39-11 |
| Sending Captures to Remote IP Addresses | 39-12 |
| Clearing Configured fcanalyzer Information | 39-12 |
| Displaying Configured Hosts | 39-13 |
| Displaying Captured Frames | 39-13 |
| Defining Display Filters | 39-14 |
| Displaying Filters Examples | 39-14 |
| Capture Filters | 39-17 |
| Permitted Capture Filters | 39-17 |
| Configuring World Wide Names | 39-18 |
| Link Initialization WWN Usage | 39-19 |
| Configuring a Secondary MAC Address | 39-19 |
| Displaying WWN Information | 39-19 |
| FC ID Allocation for HBAs | 39-20 |
| Default Company ID list | 39-20 |
| Company ID Configuration Verification | 39-21 |
| Loop Monitoring Initiation | 39-22 |
| Switch Interoperability | 39-22 |
| Configuring Interoperability | 39-24 |
| Verifying Interoperating Status | 39-25 |
| The show tech-support Command | 39-29 |
| The show tech-support brief Command | 39-30 |
| Default Settings | 39-31 |

CHAPTER 40

Configuring Fabric Configuration Servers 40-1

| | |
|-----------|------|
| About FCS | 40-1 |
|-----------|------|

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|----------------------------|------|
| Significance of FCS | 40-2 |
| FCS Name Specification | 40-2 |
| Displaying FCS Information | 40-3 |
| Default Settings | 40-6 |

CHAPTER 41

Monitoring System Processes and Logs 41-1

| | |
|---------------------------------------|-------|
| Displaying System Processes | 41-1 |
| Displaying System Status | 41-4 |
| Core and Log Files | 41-6 |
| Saving the Last Core to Flash | 41-6 |
| Clearing the Core Directory | 41-7 |
| Displaying Core Status | 41-7 |
| Kernel Core Dumps | 41-8 |
| Online System Health Management | 41-10 |
| System Health Initiation | 41-11 |
| Loopback Test Configuration Frequency | 41-11 |
| Hardware Failure Action | 41-11 |
| Test Run Requirements | 41-12 |
| Tests for a Specified Module | 41-12 |
| Clearing Previous Error Reports | 41-13 |
| Performing Internal Loopbacks | 41-14 |
| Performing External Loopbacks | 41-14 |
| Interpreting the Current Status | 41-15 |
| Displaying System Health | 41-15 |
| Default Settings | 41-18 |

INDEX

Send documentation comments to mdsfeedback-doc@cisco.com.

New and Changed Information

This chapter provides release-specific information for each new and changed feature in the Cisco MDS SAN-OS Release 2.x software. The *Cisco MDS 9000 Family Configuration Guide* is updated to address each new and changed feature in the Cisco MDS SAN-OS Release 2.x software. The latest version of this document is available at the following Cisco Systems website:

http://www.cisco.com/en/US/products/hw/ps4159/ps4358/prod_configuration_guides_list.html



Tip

The configuration guides created for previous releases are also listed in the website mentioned above. Each guide addresses the features introduced in or available in those releases. Select and view the configuration guide pertinent to the software installed in your switch.

To check for additional information about Cisco MDS SAN-OS Release 2.x, refer to the *Cisco MDS 9000 Family Release Notes* available at the following Cisco Systems website:

http://www.cisco.com/en/US/products/hw/ps4159/ps4358/prod_release_notes_list.html

Table 1 summarizes the new and changed features for the *Cisco MDS 9000 Family Configuration Guide*, and tells you where they are documented. The table includes a brief description of each new feature and the release in which the change occurred.

Table 1 **New and Changed Features for Release 2.x**

| Feature | Description | Changed in Release | Where Documented |
|--|---|--------------------|--|
| Inter-VSAN Routing (IVR) Network Address Translation (NAT) | Allows non-unique domain IDs in an IVR topology. This feature simplifies the deployment of IVR in an existing fabric. | 2.1(1a) | Chapter 17, “Configuring Inter-VSAN Routing” |
| IVR VSAN topology auto mode | Uses CFS configuration distribution in auto mode to learn the topology of the IVR-enabled switches in the network. | 2.1(1a) | Chapter 17, “Configuring Inter-VSAN Routing” |
| IVR service groups | Reduces the amount of traffic to non-IVR-enabled switches by restricting IVR-related traffic to the IVR-enabled switches. | 2.1(1a) | Chapter 17, “Configuring Inter-VSAN Routing” |
| Multiple autonomous fabric IDs (AF IDs) for IVR | Allows more than one VSAN in the network with the same VSAN ID. | 2.1(1a) | Chapter 17, “Configuring Inter-VSAN Routing” |
| IVR LUN zoning | Allows IVR to directly support LUN zoning. | 2.1(1a) | Chapter 17, “Configuring Inter-VSAN Routing” |

[Send documentation comments to mdsfeedback-doc@cisco.com.](mailto:mdsfeedback-doc@cisco.com)

Table 1 *New and Changed Features for Release 2.x (continued)*

| Feature | Description | Changed in Release | Where Documented |
|---|---|--------------------|--|
| IVZ QoS | Allows IVZ QoS to be configured separately from other zone attributes. | 2.1(1a) | Chapter 17, “Configuring Inter-VSAN Routing” |
| SANTap | Allows third-party data storage applications, such as long distance replication and continuous backup, to be integrated into the SAN. | 2.1(1a) | Chapter 25, “Configuring Intelligent Storage Services” |
| Network-Accelerated Storage Backup (NSAB) | Supports server-free backups in the SAN. | 2.1(1a) | Chapter 25, “Configuring Intelligent Storage Services” |
| Distributed configuration copy | Instructs the other switches in the fabric to save their configurations to their local NVRAM. | 2.1(1a) | Chapter 4, “Initial Configuration” |
| Enhanced IP compression auto mode | Allows auto mode option to use a combination of compression modes to effectively utilize the WAN bandwidth. | 2.1(1a) | Chapter 28, “Configuring IP Storage” |
| Zone, zone set, fcalias, and zone attribute set cloning | Allows cloning of a new zone, zone set, fcalias, or zone attribute set can be cloned from an existing zone, zone set, fcalias, or zone attribute set. | 2.1(1a) | Chapter 15, “Configuring and Managing Zones” |
| VSFN support on the Storage Services Module (SSM) | Provides support for VSFN on the SSM. | 2.1(1a) | Chapter 7, “Managing Modules” |
| File system support for log: | Allows the file system commands to support a new directory called log: for system message log files. | 2.1(1a) | Chapter 36, “Configuring System Message Logging” |
| iSCSI cut-thru routing mode | Provides iSCSI cut-thru routing mode in addition to pass-thru and store-and-forward modes. | 2.1(1a) | Chapter 28, “Configuring IP Storage” |
| New module | Storage Services Module (SSM). | 2.0(2b) | Chapter 1, “Product Overview” Chapter 7, “Managing Modules” Chapter 25, “Configuring Intelligent Storage Services” |
| Fibre Channel write acceleration | Provides support for Fibre Channel write acceleration on the ASM and SSM, which minimizes application latency or reduces transactions per second over long distances. | 2.0(2b) | Chapter 25, “Configuring Intelligent Storage Services” |
| SCSI flow statistics | Collect statistics for SCSI flows. | | |
| FICON enhancements | Provides support for FICON on MPS-14/2 modules. | 2.0(2b) | Chapter 1, “Product Overview” |
| ELP enhancement | ELP is compliant with FC-SW-3. | 2.0(2b) | Chapter 39, “Advanced Features and Concepts” |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 1 ***New and Changed Features for Release 2.x (continued)***

| Feature | Description | Changed in Release | Where Documented |
|--------------------------------------|---|--------------------|--|
| New module | 14/2-port Multiprotocol Services (MPS-14/2) module. | 2.0(1b) | Chapter 1, “Product Overview” Chapter 28, “Configuring IP Storage” Chapter 29, “Configuring IPsec Network Security” Chapter 1, “Product Overview” |
| New switches | Cisco MDS 9216i Switch. | | |
| | Cisco MDS 9216A Switch. | | |
| File system enhancements | You can use the Tab key to complete schemes, servers, and file names available in the file system | 2.0(1b) | Chapter 2, “Before You Begin” |
| Extended ping command | The ping command now provides additional options to verify the connectivity of a remote host or server. | | |
| Initial setup changes | The questions in the initial set up routine and the order in which they appear is enhanced to reflect the various changes in the Cisco SAN-OS Release 2.0(1b) software. | 2.0(1b) | Chapter 4, “Initial Configuration” |
| The show inventory command | To view information on the field replaceable units (FRUs) in the switch, including product IDs, serial numbers, and version IDs, use the show inventory command | 2.0(1b) | Chapter 8, “Managing System Hardware” |
| Cisco Fabric Services Infrastructure | The Cisco Fabric Services (CFS) infrastructure enables efficient database distribution and fosters device flexibility | 2.0(1b) | Chapter 9, “Using the CFS Infrastructure” |
| Dynamic VSANs | The Dynamic Port VSAN Membership (DPVM) feature allows you to dynamically assign VSAN membership to ports based on the device WWN. | 2.0(1b) | Chapter 11, “Creating Dynamic VSANs” |
| Graceful shutdown | The Cisco SAN-OS software implicitly performs a graceful shutdown if you shutdown an interface operating in the E port mode or if a Cisco SAN-OS software application executes a port shutdown as part of its function. | 2.0(1b) | Chapter 12, “Configuring Interfaces” |
| Extended BB_credits | To facilitate BB_credits for long haul links, the extended BB_credits flow control mechanism allows you to configure up to 3,500 receive BB_credits on a Fibre Channel port. | 2.0(1b) | |
| Small form-factor pluggable (SFP) | The term FCOT (Fibre Channel optical transmitter), is replaced by the term SFP in the Cisco SAN-OS software and in the documentation. | 2.0(1b) | |
| PortChannel | The PortChannel feature now includes a new mode (ACTIVE) and a new protocol (autocreation). | 2.0(1b) | Chapter 14, “Configuring PortChannels” |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 1 *New and Changed Features for Release 2.x (continued)*

| Feature | Description | Changed in Release | Where Documented |
|--|--|--------------------|--|
| Zone-based QoS | The zoning feature provides an additional segregation mechanism to configure the Quality of Service (QoS) priority as a zone attribute. | 2.0(1b) | Chapter 15, “Configuring and Managing Zones” Chapter 32, “Configuring Traffic Management” |
| Enhanced zoning | The zoning feature is enhanced to be compliant with FC-GS-4 and FC-SW-3. Both standards support basic zoning and enhanced zoning functionalities. | | Chapter 15, “Configuring and Managing Zones” |
| Distributed Device Alias Services | A new alias distribution feature allows you to distribute device alias names on a fabric-wide basis. | 2.0(1b) | Chapter 16, “Distributing Device Alias Services” |
| Security | Network operator default. | 2.0(1b) | Chapter 19, “Configuring Switch Security” |
| | Administrator password must be configured. | | |
| | Multiple roles support. | | |
| | Advanced Encryption Standard usage. | | |
| | Unified users and passwords. | | |
| | The error-enabled command. | | |
| The snmp-server enable traps fcdomain command | To enable a specific SNMP trap (for example, fcdomain traps) notification use the snmp-server enable traps fcdomain command. | 2.0(1b) | |
| RMON configuration | You can configure RMON alarms and events by using the CLI. | 2.0(1b) | Chapter 23, “Configuring RMON” |
| Multicast compliance | To interoperate with other vendor switches, the Cisco SAN-OS software uses the lowest domain switch as the root to compute the multicast tree in interop mode. | 2.0(1b) | Chapter 24, “Configuring Fibre Channel Routing Services and Protocols” |
| IP-ACL changes | As of Cisco SAN-OS Release 2.0(1b), you can also apply IP-ACLs to Gigabit Ethernet interfaces (IPS modules) and Ethernet PortChannel interfaces. | 2.0(1b) | Chapter 26, “Configuring IP Services” Chapter 28, “Configuring IP Storage” |
| IP storage | Tape acceleration. | 2.0(1b) | Chapter 28, “Configuring IP Storage” |
| | iSNS server. | | |
| | Mutual CHAP authentication. | | |
| | FCIP compression enhancements. | | |
| | Other changes (defaults). | | |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 1 ***New and Changed Features for Release 2.x (continued)***

| Feature | Description | Changed in Release | Where Documented |
|--------------------------------|--|---------------------------|---|
| IP Security (IPsec) | The IPsec protocol provides security services at the IP layer, including protecting one or more data flows between a pair of hosts, between a pair of security gateways, or between a security gateway and a host. | 2.0(1b) | Chapter 29, “Configuring IPsec Network Security” |
| Internet Key Exchange (IKE) | IPsec uses the IKE protocol to handle protocol and algorithm negotiation and to generate the encryption and authentication keys to be used by IPsec. | | |
| Call Home enhancements | The Call Home feature provides message throttling capabilities, periodic inventory messages, port syslog messages, and RMON alert messages. | 2.0(1b) | Chapter 30, “Configuring Call Home” |
| Port tracking | The Port Tracking feature is unique to the Cisco MDS 9000 Family. It uses information about the operational state of the link to initiate a failure in the link that connects edge device. | 2.0(1b) | Chapter 33, “Tracking and Redirecting Traffic” |
| SAN extension (SET) tuner | This feature is unique to the Cisco MDS 9000 Family. It helps you optimize FCIP performance by generating SCSI I/O commands and directing such traffic to a specific virtual target. | 2.0(1b) | Chapter 34, “Configuring the SAN Extension Tuner” |
| Command Scheduler | This feature helps you schedule configuration and maintenance jobs in any switch in the Cisco MDS 9000 Family. | 2.0(1b) | Chapter 35, “Scheduling Maintenance Jobs” |
| WWN changes | Exchange Link Protocol (ELP) and Exchange Fabric Protocol (EFP) use WWNs during link initialization. | 2.0(1b) | Chapter 39, “Advanced Features and Concepts” |
| FC ID changes | To conserve the number of FC IDs used, Cisco MDS 9000 Family switches use a special FC ID allocation scheme. | | |
| | The persistent Fibre Channel ID (FC ID) feature is enabled by default. | | Chapter 31, “Configuring Domain Parameters” Chapter 4, “Initial Configuration” |
| Storing the last core to flash | The last core dump (service core) is automatically saved to the Flash in the /mnt/pss/ partition before the switchover or reboot occurs. | 2.0(1b) | Chapter 41, “Monitoring System Processes and Logs” |

Send documentation comments to mdsfeedback-doc@cisco.com.

Preface

This preface describes the audience, organization, and conventions of the *Cisco MDS 9000 Family Configuration Guide*. It also provides information on how to obtain related documentation.

Audience

This guide is for experienced network administrators who are responsible for configuring and maintaining the Cisco MDS 9000 Family of multilayer directors and fabric switches.

Organization

This guide is organized as follows:

| Chapter | Title | Description |
|---------------------------|---|---|
| Chapter 1 | Product Overview | Presents an overview of the Cisco MDS 9000 Family of multilayer switches and directors. |
| Chapter 2 | Before You Begin | Describes the command-line interface (CLI). |
| Chapter 3 | Obtaining and Installing Licenses | Describes license types, procedure, installation, and management for the Cisco MDS SAN-OS software. |
| Chapter 4 | Initial Configuration | Provides initial switch configuration options and switch access information. |
| Chapter 5 | Configuring High Availability | Describes the high availability feature including switchover mechanisms. |
| Chapter 6 | Software Images | Describes how to upgrade Cisco MDS 9000 Family switches, install software image files, use the Flash file system on the supervisor engine, and recover a corrupted bootflash image. |
| Chapter 7 | Managing Modules | Explains how to display and analyze the status of each module and specifies the power on and power off process for modules. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| Chapter | Title | Description |
|----------------------------|---|---|
| Chapter 8 | Managing System Hardware | Explains switch hardware inventory, power usage, power supply, module temperature, fan and clock modules, and environment information. |
| Chapter 9 | Using the CFS Infrastructure | Explains the use of the Cisco Fabric Services (CFS) infrastructure to enable efficient database distribution. |
| Chapter 10 | Configuring and Managing VSANs | Describes how virtual SANs (VSANs) work, explains the concept of default VSANs, isolated VSANs, VSAN IDs, and attributes, and provides details on how to create, delete, and view VSANs. |
| Chapter 11 | Creating Dynamic VSANs | Defines the Dynamic Port VSAN Membership (DPVM) feature that is used to maintain fabric topology when a host or storage device connection is moved between two Cisco MDS switches. |
| Chapter 12 | Configuring Interfaces | Explains port and operational state concepts in Cisco MDS 9000 Family switches and provides details on configuring ports and interfaces. |
| Chapter 13 | Configuring Trunking | Explains TE ports and trunking concepts. |
| Chapter 14 | Configuring PortChannels | Explains PortChannels and load balancing concepts and provides details on configuring PortChannels, adding ports to PortChannels, and deleting ports from PortChannels. |
| Chapter 15 | Configuring and Managing Zones | Defines various zoning concepts and provides details on configuring a zone set and zone management features. |
| Chapter 16 | Distributing Device Alias Services | Describes the use of the Distributed Device Alias Services (device alias) to distribute device alias names on a fabric-wide basis. |
| Chapter 17 | Configuring Inter-VSAN Routing | Provides details on sharing resources across VSANs using the inter-VSAN Routing (IVR) feature. |
| Chapter 18 | Managing FLOGI, Name Server, FDMI, and RSCN Databases | Provides name server and fabric login details required to manage storage devices and display registered state change notification (RSCN) databases. |
| Chapter 19 | Configuring Switch Security | Discusses the AAA parameters, user profiles, RADIUS authentication, SSH services, and SNMP security options provided in all switches in the Cisco MDS 9000 Family and provides configuration information for these options. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| Chapter | Title | Description |
|----------------------------|--|---|
| Chapter 20 | Configuring Fabric Security | Describes the security protocols used in Cisco MDS switches to provide switch-switch and host-switch authentication for enterprise-wide fabrics. |
| Chapter 21 | Configuring Port Security | Provides details on port security features that can prevent unauthorized access to a switch port in the Cisco MDS 9000 Family. |
| Chapter 22 | Configuring SNMP | Provides details on configuring users, passwords, and roles for all CLI and SNMP users. |
| Chapter 23 | Configuring RMON | Provides details on using RMONs to configure alarms and events. |
| Chapter 24 | Configuring Fibre Channel Routing Services and Protocols | Provides details and configuration information on Fibre Channel routing services and protocols. |
| Chapter 25 | Configuring Intelligent Storage Services | Provides details and configuration information on the Intelligent Storage Services supported by the Storage Services Module. |
| Chapter 26 | Configuring IP Services | Provides details on IP over Fibre Channel (IPFC) services and provides configuring IPFC, virtual router, and DNS server configuration information. |
| Chapter 27 | Configuring FICON | Provides details on the FI-bre CON-nection (FICON) interface, fabric binding, and the Registered Link Incident Report (RLIR) capabilities in Cisco MDS switches. |
| Chapter 28 | Configuring IPsec Network Security | Provides details on extending the reach of Fibre Channel SANs by connecting separated SAN islands together through IP networks using FCIP, and allowing IP hosts to access FC storage using the iSCSI protocol. |
| Chapter 29 | Configuring IPsec Network Security | Provides details on the IP Security Protocol (IPsec) open standards and the Internet Key Exchange (IKE) protocol that it uses to handle protocol and algorithm negotiation. |
| Chapter 30 | Configuring Call Home | Provides details on the Call Home service and includes information on Call Home, event triggers, contact information, destination profiles, and e-mail options. |
| Chapter 31 | Configuring Domain Parameters | Explains the Fibre Channel domain (fcdomain) feature, which includes principal switch selection, domain ID distribution, FC ID allocation, and fabric reconfiguration functions. |
| Chapter 32 | Configuring Traffic Management | Provides details on the quality of service (QoS) and Fibre Channel Congestion Control (FCC) features. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| Chapter | Title | Description |
|------------|--|---|
| Chapter 33 | Tracking and Redirecting Traffic | Provides information about a port tracking feature that provides a faster recovery from link failures. |
| Chapter 34 | Configuring the SAN Extension Tuner | Explains the SAN extension tuner (SET) feature that optimizes FCIP performance. |
| Chapter 35 | Scheduling Maintenance Jobs | Describes the Cisco MDS command scheduler feature that helps you schedule configuration and maintenance jobs in any switch in the Cisco MDS 9000 Family. |
| Chapter 36 | Configuring System Message Logging | Describes how system message logging is configured and displayed. |
| Chapter 37 | Discovering SCSI Targets | Describes how the SCSI LUN discovery feature is started and displayed. |
| Chapter 38 | Monitoring Network Traffic Using SPAN | Describes the Switched Port Analyzer (SPAN), SPAN sources, filters, SPAN sessions, SD port characteristics, and configuration details. |
| Chapter 39 | Advanced Features and Concepts | Describes the advanced configuration features—time out values, fctrace, fabric analyzer, world wide names, flat FC IDs, loop monitoring, and interoperating switches. |
| Chapter 40 | Configuring Fabric Configuration Servers | Describes how the fabric configuration server (FCS) feature is configured and displayed. |
| Chapter 41 | Monitoring System Processes and Logs | Provides information on displaying system processes and status. It also provides information on configuring core and log files, HA policy, heartbeat and watchdog checks, and upgrade resets. |

Document Conventions

Command descriptions use these conventions:

| | |
|----------------------|---|
| boldface font | Commands and keywords are in boldface. |
| <i>italic font</i> | Arguments for which you supply values are in italics. |
| [] | Elements in square brackets are optional. |
| [x y z] | Optional alternative keywords are grouped in brackets and separated by vertical bars. |

Screen examples use these conventions:

| | |
|-----------------------------|---|
| screen font | Terminal sessions and information the switch displays are in screen font. |
| boldface screen font | Information you must enter is in boldface screen font. |
| <i>italic screen font</i> | Arguments for which you supply values are in italic screen font. |
| < > | Nonprinting characters, such as passwords, are in angle brackets. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|------|---|
| [] | Default responses to system prompts are in square brackets. |
| !, # | An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line. |

This document uses the following conventions:



Note

Means reader *take note*. Notes contain helpful suggestions or references to material not covered in the manual.



Caution

Means reader *be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

Related Documentation

The documentation set for the Cisco MDS 9000 Family includes the following documents:

- *Cisco MDS 9000 Family Release Notes for Cisco MDS SAN-OS Releases*
- *Cisco MDS 9000 Family Interoperability Support Matrix*
- *Cisco MDS SAN-OS Release Compatibility Matrix for IBM SAN Volume Controller Software for Cisco MDS 9000*
- *Cisco MDS SAN-OS Release Compatibility Matrix for VERITAS Storage Foundation for Networks Software*
- *Cisco MDS SAN-OS Release Compatibility Matrix for SSI Images*
- *Regulatory Compliance and Safety Information for the Cisco MDS 9000 Family*
- *Cisco MDS 9500 Series Hardware Installation Guide*
- *Cisco MDS 9200 Series Hardware Installation Guide*
- *Cisco MDS 9216 Switch Hardware Installation Guide*
- *Cisco MDS 9100 Series Hardware Installation Guide*
- *Cisco MDS 9000 Family Software Upgrade and Downgrade Guide*
- *Cisco MDS 9000 Family Configuration Guide*
- *Cisco MDS 9000 Family Command Reference*
- *Cisco MDS 9000 Family Fabric Manager Configuration Guide*
- *Cisco MDS 9000 Family Fabric and Device Manager Online Help*
- *Cisco MDS 9000 Family SAN Volume Controller Configuration Guide*
- *Cisco MDS 9000 Family MIB Quick Reference*
- *Cisco MDS 9000 Family CIM Programming Reference*
- *Cisco MDS 9000 Family System Messages Reference*
- *Cisco MDS 9000 Family Troubleshooting Guide*

Send documentation comments to mdsfeedback-doc@cisco.com.

- *Cisco MDS 9000 Family Port Analyzer Adapter 2 Installation and Configuration Note*
- *Cisco MDS 9000 Family Port Analyzer Adapter Installation and Configuration Note*

For information on VERITAS Storage Foundation™ for Networks for the Cisco MDS 9000 Family, refer to the VERITAS website: <http://support.veritas.com/>

For information on IBM TotalStorage SAN Volume Controller Storage Software for the Cisco MDS 9000 Family, refer to the IBM TotalStorage Support website:
<http://www.ibm.com/storage/support/2062-2300/>

Obtaining Documentation

Cisco documentation and additional literature are available on Cisco.com. Cisco also provides several ways to obtain technical assistance and other technical resources. These sections explain how to obtain technical information from Cisco Systems.

Cisco.com

You can access the most current Cisco documentation at this URL:

<http://www.cisco.com/univercd/home/home.htm>

You can access the Cisco website at this URL:

<http://www.cisco.com>

You can access international Cisco websites at this URL:

http://www.cisco.com/public/countries_languages.shtml

Documentation DVD

Cisco documentation and additional literature are available in a Documentation DVD package, which may have shipped with your product. The Documentation DVD is updated regularly and may be more current than printed documentation. The Documentation DVD package is available as a single unit.

Registered Cisco.com users (Cisco direct customers) can order a Cisco Documentation DVD (product number DOC-DOCDVD=) from the Ordering tool or Cisco Marketplace.

Cisco Ordering tool:

<http://www.cisco.com/en/US/partner/ordering/>

Cisco Marketplace:

<http://www.cisco.com/go/marketplace/>

Ordering Documentation

You can find instructions for ordering documentation at this URL:

http://www.cisco.com/univercd/cc/td/doc/es_inpk/pdi.htm

Send documentation comments to mdsfeedback-doc@cisco.com.

You can order Cisco documentation in these ways:

- Registered Cisco.com users (Cisco direct customers) can order Cisco product documentation from the Ordering tool:
<http://www.cisco.com/en/US/partner/ordering/>
- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco Systems Corporate Headquarters (California, USA) at 408 526-7208 or, elsewhere in North America, by calling 1 800 553-NETS (6387).

Documentation Feedback

You can send comments about technical documentation to bug-doc@cisco.com.

You can submit comments by using the response card (if present) behind the front cover of your document or by writing to the following address:

Cisco Systems
Attn: Customer Document Ordering
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Cisco Product Security Overview

Cisco provides a free online Security Vulnerability Policy portal at this URL:

http://www.cisco.com/en/US/products/products_security_vulnerability_policy.html

From this site, you can perform these tasks:

- Report security vulnerabilities in Cisco products.
- Obtain assistance with security incidents that involve Cisco products.
- Register to receive security information from Cisco.

A current list of security advisories and notices for Cisco products is available at this URL:

<http://www.cisco.com/go/psirt>

If you prefer to see advisories and notices as they are updated in real time, you can access a Product Security Incident Response Team Really Simple Syndication (PSIRT RSS) feed from this URL:

http://www.cisco.com/en/US/products/products_psirt_rss_feed.html

Reporting Security Problems in Cisco Products

Cisco is committed to delivering secure products. We test our products internally before we release them, and we strive to correct all vulnerabilities quickly. If you think that you might have identified a vulnerability in a Cisco product, contact PSIRT:

- Emergencies—security-alert@cisco.com
- Nonemergencies—psirt@cisco.com

Send documentation comments to mdsfeedback-doc@cisco.com.

**Tip**

We encourage you to use Pretty Good Privacy (PGP) or a compatible product to encrypt any sensitive information that you send to Cisco. PSIRT can work from encrypted information that is compatible with PGP versions 2.x through 8.x.

Never use a revoked or an expired encryption key. The correct public key to use in your correspondence with PSIRT is the one that has the most recent creation date in this public key server list:

<http://pgp.mit.edu:11371/pks/lookup?search=psirt%40cisco.com&op=index&exact=on>

In an emergency, you can also reach PSIRT by telephone:

- 1 877 228-7302
- 1 408 525-6532

Obtaining Technical Assistance

For all customers, partners, resellers, and distributors who hold valid Cisco service contracts, Cisco Technical Support provides 24-hour-a-day, award-winning technical assistance. The Cisco Technical Support Website on Cisco.com features extensive online support resources. In addition, Cisco Technical Assistance Center (TAC) engineers provide telephone support. If you do not hold a valid Cisco service contract, contact your reseller.

Cisco Technical Support Website

The Cisco Technical Support Website provides online documents and tools for troubleshooting and resolving technical issues with Cisco products and technologies. The website is available 24 hours a day, 365 days a year, at this URL:

<http://www.cisco.com/techsupport>

Access to all tools on the Cisco Technical Support Website requires a Cisco.com user ID and password. If you have a valid service contract but do not have a user ID or password, you can register at this URL:

<http://tools.cisco.com/RPF/register/register.do>

**Note**

Use the Cisco Product Identification (CPI) tool to locate your product serial number before submitting a web or phone request for service. You can access the CPI tool from the Cisco Technical Support Website by clicking the **Tools & Resources** link under Documentation & Tools. Choose **Cisco Product Identification Tool** from the Alphabetical Index drop-down list, or click the **Cisco Product Identification Tool** link under Alerts & RMAs. The CPI tool offers three search options: by product ID or model name; by tree view; or for certain products, by copying and pasting **show** command output. Search results show an illustration of your product with the serial number label location highlighted. Locate the serial number label on your product and record the information before placing a service call.

Send documentation comments to mdsfeedback-doc@cisco.com.

Submitting a Service Request

Using the online TAC Service Request Tool is the fastest way to open S3 and S4 service requests. (S3 and S4 service requests are those in which your network is minimally impaired or for which you require product information.) After you describe your situation, the TAC Service Request Tool provides recommended solutions. If your issue is not resolved using the recommended resources, your service request is assigned to a Cisco TAC engineer. The TAC Service Request Tool is located at this URL:

<http://www.cisco.com/techsupport/servicerequest>

For S1 or S2 service requests or if you do not have Internet access, contact the Cisco TAC by telephone. (S1 or S2 service requests are those in which your production network is down or severely degraded.) Cisco TAC engineers are assigned immediately to S1 and S2 service requests to help keep your business operations running smoothly.

To open a service request by telephone, use one of the following numbers:

Asia-Pacific: +61 2 8446 7411 (Australia: 1 800 805 227)

EMEA: +32 2 704 55 55

USA: 1 800 553-2447

For a complete list of Cisco TAC contacts, go to this URL:

<http://www.cisco.com/techsupport/contacts>

Definitions of Service Request Severity

To ensure that all service requests are reported in a standard format, Cisco has established severity definitions.

Severity 1 (S1)—Your network is “down,” or there is a critical impact to your business operations. You and Cisco will commit all necessary resources around the clock to resolve the situation.

Severity 2 (S2)—Operation of an existing network is severely degraded, or significant aspects of your business operation are negatively affected by inadequate performance of Cisco products. You and Cisco will commit full-time resources during normal business hours to resolve the situation.

Severity 3 (S3)—Operational performance of your network is impaired, but most business operations remain functional. You and Cisco will commit resources during normal business hours to restore service to satisfactory levels.

Severity 4 (S4)—You require information or assistance with Cisco product capabilities, installation, or configuration. There is little or no effect on your business operations.

Send documentation comments to mdsfeedback-doc@cisco.com.

Obtaining Additional Publications and Information

Information about Cisco products, technologies, and network solutions is available from various online and printed sources.

- Cisco Marketplace provides a variety of Cisco books, reference guides, and logo merchandise. Visit Cisco Marketplace, the company store, at this URL:

<http://www.cisco.com/go/marketplace/>

- *Cisco Press* publishes a wide range of general networking, training and certification titles. Both new and experienced users will benefit from these publications. For current Cisco Press titles and other information, go to Cisco Press at this URL:

<http://www.ciscopress.com>

- *Packet* magazine is the Cisco Systems technical user magazine for maximizing Internet and networking investments. Each quarter, Packet delivers coverage of the latest industry trends, technology breakthroughs, and Cisco products and solutions, as well as network deployment and troubleshooting tips, configuration examples, customer case studies, certification and training information, and links to scores of in-depth online resources. You can access Packet magazine at this URL:

<http://www.cisco.com/packet>

- *iQ Magazine* is the quarterly publication from Cisco Systems designed to help growing companies learn how they can use technology to increase revenue, streamline their business, and expand services. The publication identifies the challenges facing these companies and the technologies to help solve them, using real-world case studies and business strategies to help readers make sound technology investment decisions. You can access iQ Magazine at this URL:

<http://www.cisco.com/go/iqmagazine>

- *Internet Protocol Journal* is a quarterly journal published by Cisco Systems for engineering professionals involved in designing, developing, and operating public and private internets and intranets. You can access the Internet Protocol Journal at this URL:

<http://www.cisco.com/ipj>

- World-class networking training is available from Cisco. You can view current offerings at this URL:

<http://www.cisco.com/en/US/learning/index.html>



Product Overview

The Cisco MDS 9000 Family of multilayer directors and fabric switches offers intelligent fabric-switching services that realize maximum performance while ensuring high reliability levels. They combine robust and flexible hardware architecture with multiple layers of network and storage management intelligence. This powerful combination enables highly available, scalable storage networks that provide advanced security and unified management features.

The Cisco MDS 9000 Family provides intelligent networking features such as multiprotocol and multitransport integration, virtual SANs (VSANs), advanced security, sophisticated debug analysis tools, and unified SAN management.

This chapter lists the hardware features for the Cisco MDS 9000 Family and describes its software features. It includes the following sections:

- [Hardware Overview, page 1-1](#)
- [Software Features, page 1-4](#)
- [Tools for Software Configuration, page 1-14](#)

Hardware Overview

This section provides an overview of the Cisco MDS 9000 Family of multilayer directors and fabric switches.

- Cisco MDS 9120 multilayer switches contain 20 ports (4 full rate ports, 16 host-optimized ports).
- Cisco MDS 9140 multilayer switches contain 40 ports (8 full rate ports, 32 host-optimized ports).
- Cisco MDS 9216 multilayer fabric switches contain one fixed integrated supervisor module with 16 Fibre Channel ports and an expansion slot that can support up to 32 additional ports (for a total of 48 ports).
- Cisco MDS 9216A multilayer fabric switches contain one fixed integrated supervisor module with 16 Fibre Channel ports and an expansion slot that can support up to 32 additional ports (for a total of 48 ports).
- Cisco MDS 9216i multiprotocol fabric switches contain one fixed integrated supervisor module with 14 Fibre Channel ports, 2 IP ports that can support FCIP and iSCSI protocols simultaneously, and an expansion slot that can support up to 32 additional ports (for a total of 48 ports).
- Cisco MDS 9506 multilayer directors contain two slots for supervisor modules and 4 slots for switching or services modules, providing up to 128 ports (32 ports x 4 slots).
- Cisco MDS 9509 multilayer directors contain two slots for supervisor modules and 7 slots for switching or services modules, providing up to 224 ports (32 ports x 7 slots).

[Send documentation comments to mdsfeedback-doc@cisco.com.](mailto:mdsfeedback-doc@cisco.com)

Cisco MDS 9100 Series Fixed Configuration Fabric Switches

Cisco MDS 9100 Series includes two multilayer, fixed configuration (non-modular) switches:

- The Cisco MDS 9120 has 20 ports (4 full-rate ports, 16 host-optimized ports).
- The Cisco MDS 9140 provides 40 ports (8 full-rate ports, 32 host-optimized ports).

These fixed configuration switches are packaged in 1 RU enclosures and have the following features:

- Two redundant, hot-swappable power supplies have AC connections, each of which can supply power to the entire chassis.
- Two hot-swappable fan modules with two fans each manage the airflow and cooling for the entire switch.
- The 1-Gbps or 2-Gbps autosensing Fibre Channel ports support Inter-Switch Links (E ports), Extended Inter-Switch Links (TE ports), loops (FL and TL ports), and fabric (F ports) connectivity. Besides Telnet access, a 10/100BASE-T Ethernet port provides switch access.
- Hot-swappable, small form-factor pluggable (SFP) ports can be configured with either short or long wavelength SFPs for connectivity up to 500 m and 10 km, respectively.

**Note**

Switches in the Cisco MDS 9100 Series do not have a COM1 port (an RS-232 serial port).

Refer to the *Cisco MDS 9100 Series Hardware Installation Guide*.

Cisco MDS 9200 Series Fabric Switches

The Cisco MDS 9200 Series includes two multilayer switches and one multiprotocol switch:

- The Cisco MDS 9216 Switch and the Cisco MDS 9216A Switch share a consistent software architecture with the Cisco MDS 9500 Series in a semi-modular chassis and consists of the following major hardware components:
 - The chassis has two slots, one of which is reserved for the integrated supervisor module. The supervisor module provides supervisor functions and has 16 standard, Fibre Channel ports.
 - One hot-pluggable switching or services module that provides Fibre Channel or Gigabit Ethernet services.
 - The backplane has direct plug-in connectivity to one switching or services module (any type).
 - The hot-swappable fan module has four fans managing the airflow and cooling for the entire switch.
- The Cisco MDS 9216i switch shares a consistent software architecture with the Cisco MDS 9500 Series in a semi-modular chassis and consists of the following major hardware components:
 - The chassis has two slots, one of which is reserved for the integrated supervisor module. The supervisor module provides supervisor functions and has 14 standard, Fibre Channel ports and two multiprotocol ports that can support FCIP and iSCSI protocols simultaneously.
 - One hot-pluggable switching or services modules that provides Fibre Channel or Gigabit Ethernet services.
 - The backplane has direct plug-in connectivity to one switching or services module (any type).
 - The hot-swappable fan module has four fans managing the airflow and cooling for the entire switch.

Send documentation comments to mdsfeedback-doc@cisco.com.

These fabric switches have the following features:

- Two redundant, hot-swappable power supplies have AC connections, each of which can supply power to a fully loaded chassis.
- The 1-Gbps or 2-Gbps autosensing Fibre Channel ports support Inter-Switch Links (E ports), extended Inter-Switch Links (TE ports), loops (FL and TL ports), and fabric (F ports) connectivity. Besides Telnet access, a 10/100BASE-T Ethernet port provides switch access and an RS-232 (EIA/TIA-232) serial port allows switch configuration.
- Hot-swappable, small form-factor pluggable (SFP) ports can be configured with either short or long wavelength SFPs for connectivity up to 500 m and 10 km, respectively. The ports can also be configured with the extended wavelength SFPs for connectivity up to 100 km.
- The Cisco MDS 9200 Series switches support the IP Storage services (IPS) modules and the 14/2-port Multiprotocol Services (MPS-14/2) module. Both modules are configurable for both FCIP and iSCSI operation on a port-by-port basis. Ports configured for FCIP operation can be further configured to support up to three virtual ISL connections.
- The Cisco MDS 9216 switch supports the 32-port Fibre Channel Storage Services Module (SSM). The SSM enables pooling of heterogeneous storage for increased storage utilization, simplified storage management, and reduced total cost of storage ownership.

Refer to the *Cisco MDS 9216 Switch Hardware Installation Guide* and the *Cisco MDS 9200 Series Hardware Installation Guide*.

Cisco MDS 9500 Series Multilayer Directors

The Cisco MDS 9500 Series includes two multilayer, modular directors:

- The Cisco MDS 9506 Director addresses the stringent requirements of data center storage environments and consists of the following major hardware components:
 - The chassis has six slots, two of which are reserved for the supervisor modules.
 - Up to four hot-pluggable switching or services modules that provide Fibre Channel or Gigabit Ethernet services.
 - The backplane has direct plug-in connectivity to four switching or services modules, two supervisor modules, two clock modules, and two power supplies.
 - The hot-swappable fan module has six fans managing the airflow and cooling for the entire switch.
- The Cisco MDS 9509 Director addresses the stringent requirements of large data center storage environments and consists of the following major hardware components:
 - The chassis has nine slots, two of which are reserved for the supervisor modules.
 - Up to seven hot-pluggable switching or services modules that provide Fibre Channel or Gigabit Ethernet services.
 - The backplane has direct plug-in connectivity to seven switching or services modules, two supervisor modules, two clock modules, and two power supplies.
 - The hot-swappable fan module has nine fans managing the airflow and cooling for the entire switch.

Send documentation comments to mdsfeedback-doc@cisco.com.

These multilayer directors have the following features:

- Two redundant, hot-swappable power supplies have AC or DC connection, each of which can supply power to the entire chassis.
- Two supervisor modules ensure high availability and traffic load balancing capabilities. Each supervisor module can control the entire switch. The standby supervisor module provides redundancy in case the active supervisor module fails.
- The 1-Gbps or 2-Gbps autosensing Fibre Channel ports support Inter-Switch Links (E ports), Extended Inter-Switch Links (TE ports), loops (FL and TL ports), and fabric (F ports) connectivity. Besides Telnet access, a 10/100BASE-T Ethernet port provides switch access and an RS-232 serial port allows switch configuration.
- Hot-swappable, small form-factor pluggable (SFP) ports can be configured with either short or long wavelength SFPs for connectivity up to 500 m and 10 km, respectively.
- The Cisco MDS 9500 Series switches support the IP Storage Services (IPS) module and the 14/2-port Multiprotocol Services (MPS-14/2) module. Both modules are configurable for both FCIP and iSCSI operation on a port-by-port basis. Ports configured for FCIP operation can be further configured to support up to three virtual ISL connections. FICON is supported on the MPS-14/2 module and functions with all the IPS features on this module, including IPsec and hardware compression.
- The Cisco MDS 9500 Series switches support the 32-port Fibre Channel Storage Services Module (SSM). The SSM enables pooling of heterogeneous storage for increased storage utilization, simplified storage management, and reduced total cost of storage ownership.

Refer to the *Cisco MDS 9500 Series Hardware Installation Guide*.

Software Features

This section provides an overview of the major software features of the Cisco MDS 9000 Family of multilayer directors and fabric switches.

Licensing

The licensing functionality is available in all switches in the Cisco MDS 9000 Family. This functionality allows you to access specified premium features on the switch after you install the appropriate license for that feature. Licenses are sold, supported, and enforced as of Cisco MDS SAN-OS Release 1.3(1).

See [Chapter 3, “Obtaining and Installing Licenses.”](#)

High Availability

The Cisco MDS 9500 Series supports application restartability and nondisruptive supervisor switchability. The switches are protected from system failure by redundant hardware components and a high availability software framework. The high availability (HA) software framework includes the following:

- Provides stateful redundancy for supervisor module failure by using dual supervisor modules.
- Ensures nondisruptive software upgrade capability.

Send documentation comments to mdsfeedback-doc@cisco.com.

- Protects against link failure using the PortChannel (port aggregation) feature. This feature is also available in Cisco MDS 9200 Series and in the Cisco MDS 9100 Series.
- Provides management redundancy using the Virtual Router Redundancy Protocol (VRRP). This feature is also available in Cisco MDS 9200 Series and in the Cisco MDS 9100 Series.
- Performs nondisruptive restarts of a failed process on the same supervisor module. A service running on the supervisor modules and on the switching or services module tracks the HA policy defined in the configuration and takes action based on this policy. This feature is also available in the Cisco MDS 9200 Series and in the Cisco MDS 9100 Series.

See [Chapter 5, “Configuring High Availability,”](#) [Chapter 6, “Software Images,”](#) [Chapter 14, “Configuring PortChannels,”](#) and the [“The Virtual Router Redundancy Protocol”](#) section on page 26-19.

Switch Reliability

Switches in the Cisco MDS 9000 Family maintain internally controlled reliability services that ensure continued service with no degradation. This reliability service includes the following:

- Provides power-on self testing (POST)
- Detects errors, isolates faults, performs parity checking, and checks illegal addresses
- Enables remote diagnostics using Call Home troubleshooting features
- Displays LEDs that summarize the status of each switching or services module, supervisor module, power supply, and fan assembly

Graceful Shut Down

As of Release 2.0(1b), the Cisco SAN-OS software implicitly performs a graceful shut down in response to either of the following actions:

- If you shut down an interface operating in the E port mode
- If a Cisco SAN-OS software application executes a port shut down as part of its function

A graceful shut down ensures that no frames are lost when the interface is shutting down. When a shut down is triggered either by you or the Cisco SAN-OS software, the switches connected to the shut down link coordinate with each other to ensure that all frames in the ports are safely sent through the link before shutting down. This enhancement reduces the chance of frame loss.

Cisco Fabric Services

As of Release 2.0(1b), the Cisco SAN-OS software uses the Cisco Fabric Services (CFS) infrastructure to enable efficient database distribution and to foster device flexibility. It simplifies SAN provisioning by automatically distributing configuration information to all switches in a fabric. The following Cisco SAN-OS features use the CFS infrastructure:

- NTP (see [“NTP Configuration Distribution”](#) section on page 4-19).
- Dynamic Port VSAN Membership (see [Chapter 11, “Creating Dynamic VSANs”](#)).
- Distributed Device Alias Services (see [Chapter 16, “Distributing Device Alias Services”](#)).
- IVR topology (see [“Database Merge Guidelines”](#) section on page 17-23).
- TACACS and RADIUS (see the [“Distributing AAA Server Configuration”](#) section on page 19-15).

Send documentation comments to mdsfeedback-doc@cisco.com.

- User and administrator roles (see the “[Role-Based Authorization](#)” section on page 19-21).
- Port security (see the “[Port Security Configuration Distribution](#)” section on page 21-9).
- iSNS (see the “[Configuring iSCSI Storage Name Services](#)” section on page 28-106).
- Call Home (see the “[Call Home Configuration Distribution](#)” section on page 30-12).
- Syslog (see the “[System Message Logging Configuration Distribution](#)” section on page 36-7).
- Fctimer (see the “[fctimer Distribution](#)” section on page 39-3).
- SCSI Flow Services (see the “[Configuring SCSI Flow Services](#)” section on page 25-3).
- Saving the configuration (see the “[Saving the Configuration](#)” section on page 4-27).

Virtual SANs

VSANs (virtual SANs) enable higher security and greater scalability in Fibre Channel fabrics. VSANs provide isolation among devices that are physically connected to the same fabric. VSANs allow multiple logical SANs over a common physical infrastructure. VSANs offer the following:

- Traffic isolation—Traffic is contained within VSAN boundaries and devices reside only in one VSAN thus ensuring absolute separation between user groups, if desired.
- Scalability—VSANs are overlaid on top of a single physical SAN. The ability to create several logical VSAN layers increases the scalability of the SAN.
- Per VSAN fabric services—Replication of fabric services on a per VSAN basis provides increased scalability and availability.
- Redundancy—Several VSANs created on the same physical SAN ensure redundancy. If one VSAN fails, there is a configured backup path between the host and the switch.
- Ease of configuration—Devices can be added, moved, or changed between VSANs without changing the physical structure of a SAN. Moving a device from one VSAN to another only requires configuration at the port level, not at a physical level.

See [Chapter 10, “Configuring and Managing VSANs.”](#)

Dynamic VSANs

Port VSAN membership on the switch is assigned on a port-by-port basis. By default each port belongs to the default VSAN.

As of Cisco SAN-OS Release 2.0(1b), you can dynamically assign VSAN membership to ports by assigning VSANs based on the device WWN. This method is referred to as the Dynamic Port VSAN Membership (DPVM) feature. DPVM offers flexibility and eliminates the need to reconfigure the VSAN to maintain fabric topology when a host or storage device connection is moved between two Cisco MDS switches. It retains the configured VSAN regardless of where a device is connected or moved.

See [Chapter 11, “Creating Dynamic VSANs.”](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

Intelligent Zoning

Zoning controls access between devices in a VSAN. Zoning accomplishes the following:

- Partitions devices that use different operating systems. In a heterogeneous environment, it is often necessary to separate servers and storage devices to avoid accidental transfer of information between devices with different operating systems. Such transfers could result in corruption or deletion of data.
- Creates logical subsets of closed user groups. Closed user groups are needed to enforce security or to separate functional areas across the fabric.
- Configures groups of devices that are separate from the rest of the fabric. Based on the assigned zone membership, devices outside the zone cannot access devices internal to the zone.
- Provides temporary access between devices (zone sets). Zone restrictions can be imposed temporarily, and then restored to revert to normal operation, if desired.
- Restricts access to specific logical unit numbers (LUNs) associated with a device.
- Allows members to have only read-only access to the media within a read-only Fibre Channel zone.

See [Chapter 15, “Configuring and Managing Zones.”](#)

Enhanced Zoning

As of Cisco SAN-OS Release 2.0(1b), the zoning feature is enhanced to be compliant with FC-GS-4 and FC-SW-3. Both standards support the basic zoning features explained in the previous section and the enhanced zoning features discussed in this section.

See the [“About Enhanced Zoning”](#) section on page 15-27.

Device Alias Distribution

As of Release 2.0(1b), all switches in the Cisco MDS 9000 Family offer a new alias distribution feature called Distributed Device Alias Services (device alias). In Release 1.3 and earlier, aliases were distributed on a per VSAN basis. Using this new, enhanced service, you now have the option to distribute device alias names on a fabric-wide basis.

See [Chapter 16, “Distributing Device Alias Services.”](#)

Inter-VSAN Routing

Using Inter-VSAN Routing (IVR), resources across VSANs can be accessed without compromising other VSAN benefits. Valuable resources such as tape libraries are easily shared across VSANs without compromise. Routes that traverse one or more VSANs across multiple switches can be established, if necessary, to establish proper interconnections. IVR used in conjunction with FCIP provides more efficient business continuity or disaster recovery solutions.

See [Chapter 17, “Configuring Inter-VSAN Routing.”](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

Trunking

Trunking is the term used to refer to an ISL link that carries one or more VSANs. Trunking ports receive and transmit Enhanced ISL (EISL) frames. EISL frames carry an EISL header containing VSAN information. Once EISL is enabled on an E port, that port becomes a TE port. The trunking configuration is saved along with the interface information.

See [Chapter 12, “Configuring Interfaces”](#) and [Chapter 13, “Configuring Trunking.”](#)

PortChannels

PortChannel refers to the aggregation of multiple physical Fibre Channel ports into one logical port to provide high aggregated bandwidth, load balancing, and link redundancy. Up to 16 physical ports can be aggregated into a PortChannel. PortChannels can connect to ports across switching or services modules. The failure of a port in one module does not bring down the logical PortChannel link. Specifically, a PortChannel does the following:

- Increases the aggregated bandwidth on an ISL or EISL by distributing traffic among all functional links in the channel.
- Load balances across multiple links and maintains optimum bandwidth utilization. Load balancing is based on a source ID (SID), a destination ID (DID), and optionally an originator exchange ID (OX ID) which identify the flow of the frame.
- Provides high availability on an ISL. If one link fails, traffic previously carried on this link is switched to the remaining links. If a link goes down in a PortChannel, the upper protocol is not aware of it. To the upper protocol, the link is still there, although the bandwidth is diminished. The routing tables are not affected by link failure. PortChannels can contain up to 16 physical links and can span multiple modules for added high availability.
- As of Cisco SAN-OS Release 2.0(1b), a protocol to exchange PortChannel configurations is available in all Cisco MDS switches. This addition simplifies PortChannel management with incompatible ISLs. An additional autocreation mode enables ISLs with compatible parameters to automatically form channel groups without manual intervention.

See [Chapter 14, “Configuring PortChannels.”](#)

IP Services

Switches in the Cisco MDS 9000 Family support the following IP services:

- IP over Ethernet —These services are limited to management traffic.
- IP over Fibre Channel (IPFC)—IPFC (RFC 2625) specifies how IP packets are transported using encapsulation schemes. By encapsulating IP frames into Fibre Channel frames, management information is exchanged among switches without requiring a separate Ethernet connection to each switch. Each switch includes:
 - Encapsulation for IP and the Address Resolution Protocol (ARP) over Fibre Channel.
 - Address resolution uses the ARP server.
- IP routing services—These services include:
 - Ethernet or TCP/IP connection.
 - Static IP routing services to enable management traffic between VSANs.

Send documentation comments to mdsfeedback-doc@cisco.com.

- DNS client support.
- The Network Time Protocol (NTP) server to synchronize the system clocks of network devices.
- In Cisco SAN-OS Release 1.3 and earlier, you could only apply IP-ACLs to VSAN interfaces and the management interface. As of Cisco SAN-OS Release 2.0(1b), you can also apply IP-ACLs to Gigabit Ethernet interfaces (IPS modules) and Ethernet PortChannel interfaces.

See [Chapter 26, “Configuring IP Services.”](#)

FICON

Fibre Connection (FICON) interface capabilities enhance the Cisco MDS 9000 Family by supporting both open systems and mainframe storage network environments. Inclusion of Control Unit Port (CUP) support further enhances the MDS offering by allowing in-band management of the switch from FICON processors.

See the [Chapter 27, “Configuring FICON.”](#)

Fabric Binding

The fabric binding feature helps prevent unauthorized switches from joining the fabric or disrupting current fabric operations.

See the [“Fabric Binding Configuration” section on page 27-37.](#)

RLIR

The Registered Link Incident Report ((RLIR) application provides a method for a switchport to send a LIR to a registered Nx-port.

See the [“Displaying RLIR Information” section on page 27-44.](#)

IP Storage

The Cisco MDS 9000 Family IP services module, the 14/2-port Multiprotocol Service module, and the Cisco MDS 93126i Switch integrate seamlessly into the Cisco MDS 9000 Family of multilayer directors and fabric switches. Traffic can be routed between any IP storage port and any other port on a Cisco MDS 9000 Family switch. These products support the full range of services available on other Cisco MDS 9000 Family switching modules including VSANs, security, and traffic management. It uses widely known IP to cost-effectively connect to more servers and more locations over greater distances than previously possible. It delivers both Fibre Channel over IP (FCIP) and iSCSI IP storage services and is configurable on a port-by-port basis.

- FCIP highlights
 - Simplifies data protection and business continuance strategies by enabling backup, remote replication, and disaster recovery over WAN distances using open-standard FCIP tunneling.
 - Improves utilization of WAN resources for backup and replication by tunneling up to three virtual Inter-Switch Links (ISLs) on a single Gigabit Ethernet port.
 - Reduces SAN complexity by eliminating the need to deploy and manage a separate remote connectivity platform.

Send documentation comments to mdsfeedback-doc@cisco.com.

- Preserves Cisco MDS 9000 Family enhanced capabilities including VSANs, advanced traffic management, and security across remote connections.
- Improves application performance using one or more of the following options for the FCIP interface: FCIP write acceleration, FCIP tape acceleration, and FCIP compression.
- iSCSI highlights
 - Extends the benefits of Fibre Channel SAN-based storage to IP-enabled servers at a lower cost point than possible using Fibre Channel interconnect alone.
 - Increases storage utilization and availability through consolidation of IP and Fibre Channel block storage.
 - Preserves through a transparent operation the functionality of legacy storage applications such as zoning tools.
 - Allows your existing TCP/IP networks to function more effectively as storage area networks by automating the discovery, management, and configuration of iSCSI devices.

See [Chapter 28, “Configuring IP Storage.”](#)

Call Home

The Call Home feature detects switch failures and sends alerts along with relevant failure information. These alerts are sent through e-mail to a user-specified customer center.

As of Cisco SAN-OS Release 2.0(1b), the Call Home feature provides message throttling capabilities, periodic inventory messages, port syslog messages, and RMON alert messages.

See [Chapter 30, “Configuring Call Home.”](#)

QoS and Congestion Control

Switches in the Cisco MDS 9000 Family provide priority queuing and flow control services.

- The Quality of Service (QoS) feature has the following advantages:
 - Guarantees relative bandwidth to application traffic.
 - Controls latency experienced by application traffic.
 - Prioritizes one application over another (for example, prioritizing transactional traffic over bulk traffic) through bandwidth and latency differentiation.
- Fibre Channel Congestion Control (FCC)—FCC is a flow control mechanism that alleviates congestion on Fibre Channel networks. Any switch in the network can detect congestion for an output port. The switches sample frames from the congested queue and generate messages about the congestion level upstream toward the source of the congestion. The switch closest to the source, with FCC enabled, can perform one of two actions:
 - Forwards the frames as other vendor switches do.
 - Limits the flow of frames from the port causing the congestion.

See [Chapter 32, “Configuring Traffic Management.”](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

SPAN and RSPAN

The Switched Port Analyzer (SPAN) feature is specific to switches in the Cisco MDS 9000 Family. It monitors network traffic through a Fibre Channel interface. Traffic through any Fibre Channel interface can be replicated to a special port called the SPAN destination port (SD port). Any Fibre Channel port in a switch can be configured as an SD port. Once an interface is in SD-port mode, it cannot be used for normal data traffic. You can attach a Fibre Channel Analyzer to the SD port to monitor SPAN traffic.

The Remote SPAN (RSPAN) feature enables you to remotely monitor traffic for one or more SPAN sources distributed in one or more source switches in a Fibre Channel fabric. The SPAN destination (SD) port is used for remote monitoring in a destination switch. A destination switch may be different from the source switch(es) provided that it is attached to the same Fibre Channel fabric. You can replicate and monitor traffic in any remote Cisco MDS 9000 Family switch or director, just as you would monitor traffic in an MDS source switch. This feature is nonintrusive and does not affect network traffic switching for any SPAN source ports.

See [Chapter 38, “Monitoring Network Traffic Using SPAN.”](#)

Switch Management Features

Besides the software features already listed, there are additional management features that fall into the following categories: redundant supervisor module management, fabric management, and security management.

Redundant Supervisor Module Management

The Cisco MDS 9500 Series of multilayer directors support two redundant supervisor modules. They require two supervisor modules to enforce redundant supervisor module management and high availability (see [Table 1-1](#)).

Table 1-1 Supervisor Module Options in Cisco MDS 9000 Switches

| Product | No. of Supervisor Modules | Supervisor Module Slot No. | Switching/Services Module Features |
|-----------------------|---|----------------------------|--|
| Cisco MDS 9100 Series | Not applicable | | |
| Cisco MDS 9200 Series | One module (includes 16 additional ports) | 1 | 2-slot chassis allows one optional switching or services module in the other slot. |
| Cisco MDS 9506 | Two modules | 5 and 6 | 6-slot chassis allows any switching or services module in the other four slots. |
| Cisco MDS 9509 | Two modules | 5 and 6 | 9-slot chassis allows any switching or services module in the other seven slots. |

When a switch powers up and two supervisor modules are present, the module in slot 5 enters the active mode, while the second module in slot 6 enters the standby mode. All storage management functions occur on the active supervisor module. The standby module constantly monitors the active module. If the active module fails, the standby module takes over without any impact to user traffic.

Refer to the *Cisco MDS 9500 Series Hardware Installation Guide*.

Send documentation comments to mdsfeedback-doc@cisco.com.

Fabric Management

Switches in the Cisco MDS 9000 Family offer fabric management and control through the command-line interface (CLI) by using Telnet, SSH, or a serial console and through the Cisco MDS 9000 Fabric Manager tool by using the Simple Network Management Protocol (SNMP) services:

- SNMP versions 1, 2, and 3 are supported. See [Chapter 19, “Configuring Switch Security.”](#)
- Remote Monitoring (RMON) allows you to specify thresholds and monitor alarms on SNMP variables. Extended RMON alarms are available for supported Management Information Base (MIB) objects (refer to the *Cisco MDS 9000 Family MIB Reference*). See [Chapter 23, “Configuring RMON.”](#)
- System log (syslog) messages are viewed through a console or Telnet session for asynchronous events such as an interface transition. System messages are directed to an internal log and optionally to an external server (refer to the *Cisco MDS 9000 Family System Messages Reference*). See [Chapter 36, “Configuring System Message Logging.”](#)

Security Management

The Cisco MDS 9000 Family of switches offer strict and secure switch management options through switch access security, port security, user authentication, and role-based access control.

Network Security

IP Security Protocol (IPsec) is a framework of open standards that provides data confidentiality, data integrity, and data authentication between participating peers. It is developed by the Internet Engineering Task Force (IETF). IPsec provides these security services at the IP layer. IPsec can be used to protect one or more data flows between a pair of hosts, between a pair of security gateways, or between a security gateway and a host.

IPsec uses the Internet Key Exchange (IKE) protocol to handle protocol and algorithm negotiation and to generate the encryption and authentication keys to be used by IPsec. While IKE can be used with other protocols, its initial implementation is with the IPsec protocol. IKE provides authentication of the IPsec peers, negotiates IPsec security associations, and establishes IPsec keys.

See [Chapter 29, “Configuring IPsec Network Security.”](#)

Fabric Security

Fibre Channel Security Protocol (FC-SP) capabilities in Cisco MDS SAN-OS Release 1.3 provides switch-switch and host-switch authentication to overcome security challenges for enterprise-wide fabrics. Diffie-Hellman Challenge Handshake Authentication Protocol (DHCHAP) is an FC-SP protocol implemented in this release to provide authentication between Cisco MDS 9000 Family switches and other devices. It consists of the CHAP protocol combined with the Diffie-Hellman exchange.

See [Chapter 20, “Configuring Fabric Security.”](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

Switch Access Security

Each switch can be accessed through the CLI or SNMP.

- Secure switch access—Available when you explicitly enable Secure Shell Protocol (SSH) access to the switch. SSH access provides additional controlled security by encrypting data, user IDs, and passwords. By default, Telnet access is enabled on each switch.
- SNMP access—SNMPv3 provides built-in security for secure user authentication and data encryption.
- IP access control lists (IP-ACLs)—IP-ACLs provide basic network security to all switches in the Cisco MDS 9000 Family. IP-ACLs restrict IP-related inband and out-of-band management traffic based on IP addresses (Layer 3 and Layer 4 information). You can use IP-ACLs to control transmissions on an interface.

See [Chapter 19, “Configuring Switch Security.”](#)

Port Security

Port security features prevent unauthorized access to a switch port in the Cisco MDS 9000 Family.

- Login requests from unauthorized Fibre Channel devices (Nx ports) and switches (xE ports) are rejected.
- All intrusion attempts are reported to the SAN administrator through system messages.

See [Chapter 21, “Configuring Port Security.”](#)

User Authentication

A strategy known as authentication, authorization, and accounting (AAA) is used to verify identity of, grant access, and track the actions of remote users. The Remote Access Dial-In User Service (RADIUS) and Terminal Access Controller Access Control System Plus (TACACS+) provide AAA solutions.

Based on the user ID and password combination provided, switches perform local authentication using a local database or remote authentication using AAA server(s). A global, preshared, secret key authenticates communication between the AAA servers. This secret key can be configured for all AAA server groups or for only a specific AAA server. This kind of authentication provides a central configuration management capability.

See [Chapter 19, “Configuring Switch Security.”](#)

Role-Based Access

Role-based access control assigns roles or groups (locally through the switch or remotely using AAA servers) to users and limits access to the switch. Access is assigned based on the permission level associated with each user ID. Your administrator can provide complete access to each user or restrict access to specific read and write levels for each command.

As of Cisco MDS SAN-OS Release 1.2(x), CLI and SNMP in all switches in the Cisco MDS 9000 Family synchronize CLI and SNMP roles. You can use SNMP to modify a role that was created using CLI and vice versa. Each role in SNMP is the same as a role created or modified through the CLI.

Each role is restricted to one or more VSAN as required.

See [Chapter 19, “Configuring Switch Security.”](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

Port Tracking

The Port Tracking feature is unique to the Cisco MDS 9000 Family of switches. This feature uses information about the operational state of the link to initiate a failure in the link that connects the edge device. This process of converting the indirect failure to a direct failure triggers a faster recovery process towards redundant links. When enabled, the port tracking feature brings down the configured links based on the failed link and forces the traffic to be redirected to another redundant link.

See [Chapter 33, “Tracking and Redirecting Traffic.”](#)

SAN Extension Tuner

The SAN extension tuner (SET) feature is unique to the Cisco MDS 9000 Family of switches. This feature helps you optimize FCIP performance by generating SCSI I/O commands and directing such traffic to a specific virtual target. You can specify the size of the test I/O transfers and how many concurrent I/Os to generate while testing. The SET reports the resulting I/Os per second (IOPS) and I/O latency, which helps you determine the number of concurrent I/Os needed to maximize FCIP throughput.

See [Chapter 34, “Configuring the SAN Extension Tuner.”](#)

Command Scheduler

The Cisco MDS command scheduler feature helps you schedule configuration and maintenance jobs in any switch in the Cisco MDS 9000 Family. This feature is available in the Cisco SAN-OS Release 2.0(1b) software and later. You can use this feature to schedule jobs on a one-time basis or periodically.

See [Chapter 35, “Scheduling Maintenance Jobs.”](#)

Intelligent Storage Services

The Advanced Services Modules (ASMs) and Storage Services Modules (SSMs) support Intelligent Storage Services. Intelligent Storage Services include the following:

- Fibre Channel write acceleration
- SCSI flow statistics

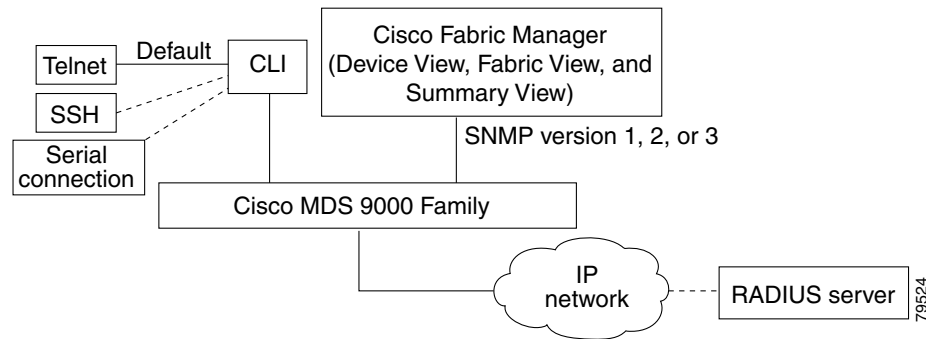
See [Chapter 25, “Configuring Intelligent Storage Services.”](#)

Tools for Software Configuration

You can use one of two configuration management tools to configure your SANs: the CLI and the Cisco MDS 9000 Fabric Manager graphical user interface (see [Figure 1-1](#)).

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 1-1 Tools for Configuring Software



CLI

With the CLI, you can type commands at the switch prompt, and the commands are executed when you press the Enter key. The CLI parser provides command help, command completion, and keyboard sequences that allow you to access previously executed commands from the buffer history.

Continue reading this guide for more information on configuring the Cisco MDS switch using the CLI.

Cisco MDS 9000 Fabric Manager

The Cisco Fabric Manager is a set of network management tools that supports Secure Simple Network Management Protocol version 3 (SNMPv3) and legacy versions. It provides a graphical user interface (GUI) that displays real-time views of your network fabric, and lets you manage the configuration of Cisco MDS 9000 Family devices and third-party switches. The Cisco Fabric Manager applications are:

- Fabric Manager Server—performs advanced monitoring, troubleshooting, and configuration for multiple fabrics. It must be started before running the Fabric Manager. It can be accessed by up to 16 Fabric Manager clients at a time.
- Device Manager—presents two views of a switch.
 - Device View displays a continuously updated physical representation of the switch configuration, and provides access to statistics and configuration information for a single switch.
 - Summary View presents real-time performance statistics of all active interfaces and channels on the switch for Fibre Channel and IP connections.
- Fabric Manager Web Client—allows operators to monitor MDS events, performance, and inventory from a remote location using a web browser.
- Performance Manager—provides detailed traffic analysis by capturing data with SNMP. This data is compiled into various graphs and charts which can be viewed with any web browser.

The Cisco Fabric Manager applications are an alternative to the CLI for most switch configuration commands.



Note

Resource Manager Essentials (RME) versions 3.4 and 3.5 provide support for switches in the Cisco MDS 9000 Family. Device Updates (DU) are available on Cisco.com (<http://www.cisco.com/>).

Refer to the *Cisco MDS 9000 Fabric Manager Configuration Guide*.

Send documentation comments to mdsfeedback-doc@cisco.com.



Before You Begin

This chapter prepares you to configure switches from the CLI. It also lists the information you need to have before you begin, and it describes the CLI command modes.

This chapter includes the following sections:

- [About the Switch Prompt, page 2-2](#)
- [Default Switch Roles, page 2-3](#)
- [About the CLI Command Modes, page 2-3](#)
- [CLI Command Hierarchy, page 2-4](#)
- [CLI Command Navigation, page 2-8](#)
- [About Flash Devices, page 2-20](#)
- [Formatting Flash Devices and File Systems, page 2-21](#)
- [Using the File System, page 2-22](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

About the Switch Prompt



Note

Refer to the *Cisco MDS 9200 Series Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide* for installation and connection instructions.

Once the switch is powered on successfully, you see the default switch prompt (switch#) as shown in [Example 2-1](#).

Example 2-1 Output When Switch Boots Up

```
Auto booting bootflash:/boot-279 bootflash:/system_image;...
Booting kickstart image:bootflash:/boot-279....
.....Image verification OK

Starting kernel...
INIT: version 2.78 booting
Checking all filesystems..... done.
Loading system software
Uncompressing system image: bootflash:/system_image
CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
INIT: Entering runlevel: 3

<<<<<SAN OS bootup log messages>>>>>

      ---- Basic System Configuration Dialog ----

This setup utility will guide you through the basic configuration of
the system. Use ctrl-c to abort configuration dialog at any prompt.

Basic management setup configures only enough connectivity for
management of the system.

Would you like to enter the basic configuration dialog (yes/no): yes

<<<<<after configuration>>>>>

switch login:admin101
Password:*****
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2004, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
Cisco Systems, Inc. and/or other third parties and are used and
distributed under license. Some parts of this software are covered
under the GNU Public License. A copy of the license is available
at http://www.gnu.org/licenses/gpl.html.
switch#
```

You can perform embedded CLI operations, access command history, and use command parsing functions at this prompt. The switch gathers the command string upon detecting an **Enter** (CR) and accepts commands from a terminal.

Send documentation comments to mdsfeedback-doc@cisco.com.

Default Switch Roles

By default, two roles exist in all switches:

- Network operator—Has permission to view the configuration.
- Network administrator—Has permission to perform all functions and to set up to 64 permission levels based on user roles and groups.

When you execute a command, perform command completion, or obtain context sensitive help, the switch software allows the operation to progress if you have the correct permission as specified in the description of the command. see [Chapter 19, “Configuring Switch Security.”](#)

About the CLI Command Modes

Switches in the Cisco MDS 9000 Family have two main command modes—user EXEC mode and configuration mode. The commands available to you depend on the mode you are in. To obtain a list of available commands in either mode, type a question mark (?) at the system prompt.

[Table 2-1](#) lists and describes the two commonly used modes, how to enter the modes, and the resulting system prompts. The system prompt helps you identify which mode you are in and hence, which commands are available to you.

Table 2-1 *Frequently Used Switch Command Modes*

| Mode | Description of Use | How to Access | Prompt |
|--------------------|---|---|-----------------|
| EXEC | Enables you to temporarily change terminal settings, perform basic tests, and display system information. Note Changes made in this mode are generally not saved across system resets. | At the switch prompt, enter the required EXEC mode command. | switch# |
| Configuration mode | Enables you to configure features that affect the system as a whole. Note Changes made in this mode are saved across system resets if you save your configuration. See the “Saving a Configuration” section on page 2-13. | From EXEC mode, enter the config terminal command. | switch(config)# |

You can abbreviate commands and keywords by entering just enough characters to make the command unique from other commands. For example, you can abbreviate the **config terminal** command to **conf t**.



Note

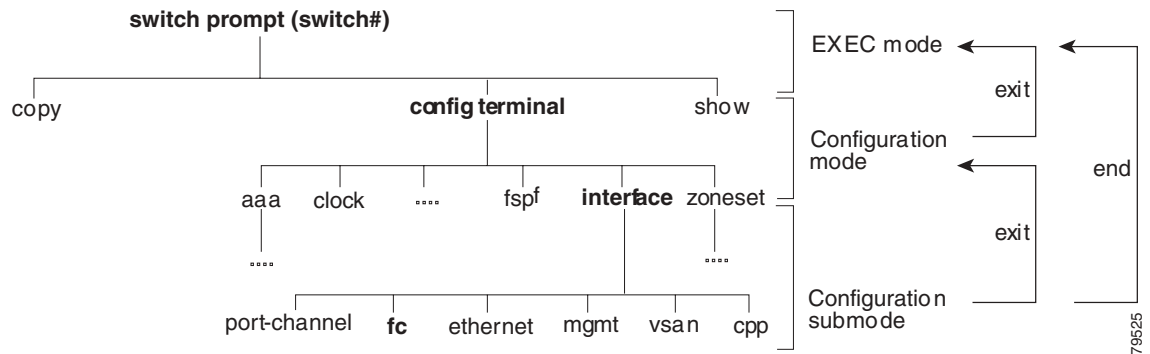
Do not enter percent (%), pound (#), ellipsis (...), vertical bar (|), less than or great than (< >), brackets ([]), or braces ({ }) in command lines. These characters have special meaning in Cisco SAN-OS text strings.

Send documentation comments to mdsfeedback-doc@cisco.com.

CLI Command Hierarchy

The CLI commands are organized hierarchically, with commands that perform similar functions grouped under the same level. For example, all commands that display information about the system, configuration, or hardware are grouped under the **show** command, and all commands that allow you to configure the switch are grouped under the **config terminal** command. Figure 2-1 illustrates a portion of the **config terminal** command hierarchy.

Figure 2-1 CLI Command Hierarchy Example



To execute a command, you enter the command by starting at the top level of the hierarchy. For example, to configure a Fibre Channel interface, use the **config terminal** command. Once you are in configuration mode, issue the **interface** command. When you are in the interface submenu, you can query the available commands there.

The following example shows how to query the available commands in the interface submenu:

```

switch# config t
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# interface fc1/1
switch(config-if)# ?
Interface configuration commands:
  channel-group  Add to/remove from a port-channel
  do             EXEC command
  exit           Exit from this submenu
  fcdomain       Configure fcdomain parameters
  fspf           Configure FSPF parameters
  no             Negate a command or set its defaults
  rspan-tunnel   Configure remote span tunnel interface
  shutdown       Enable/disable an interface
  switchport     Configure switchport parameters
  
```

Send documentation comments to mdsfeedback-doc@cisco.com.

EXEC Mode Options

When you start a session on the switch, you begin in EXEC mode. Based on the role or group to which you belong, you have access to limited commands or to all commands (see the [“Role-Based Authorization” section on page 19-21](#)). From EXEC mode, you can enter configuration mode. Most of the EXEC commands are one-time commands, such as **show** commands, which display the current configuration status. Here is a list of EXEC mode commands:

```
switch# ?
Exec commands:
  attach      Connect to a specific linecard
  callhome    Callhome commands
  cd          Change current directory
  clear       Reset functions
  clock       Manage the system clock
  config      Enter configuration mode
  copy        Copy from one file to another
  crypto      Act on crypto associations
  debug       Debugging functions
  delete      Delete a file
  dir         List files in a directory
  discover    Discover information
  exit        Exit from the EXEC
  fcping      Ping an N-Port
  fctrace     Trace the route for an N-Port.
  find        Find a file below the current directory
  format      Format disks
  gunzip      Uncompresses LZ77 coded files
  gzip        Compresses file using LZ77 coding
  install     Upgrade software
  ips         Various sbyte module related commands
  ivr         IVR exec commands
  mkdir       Create new directory
  modem       Modem commands
  move        Move files
  no          Disable debugging functions
  ping        Send echo messages
  port-channel Port-Channel related commands
  purge       Deletes unused data
  pwd         View current directory
  reload      Reboot the entire box
  rmdir       Delete a directory
  run-script  Run shell scripts
  send        Send message to open sessions
  setup       Run the basic SETUP command facility
  show        Show running system information
  sleep       Sleep for the specified number of seconds
  ssh         SSH to another system
  system      System management commands
  tac-pac     Save tac information to a specific location
  tail        Display the last part of a file
  telnet      Telnet to another system
  terminal    Set terminal line parameters
  test        Test command
  traceroute  Trace route to destination
  undebg      Disable Debugging functions (See also debug)
  update      Update license
  write       Write current configuration
  zone        Execute Zone Server commands
  zoneset     Execute zoneset commands
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Configuration Mode

In configuration mode, you can make changes to the existing configuration. When you save the configuration, these commands are preserved across switch reboots. Once you are in configuration mode, you can enter interface configuration mode, zone configuration mode, and a variety of protocol-specific modes. Configuration mode is the starting point for all configuration commands. When you are in configuration mode, the switch expects configuration commands from the user.

The following example shows output from the **config terminal** command:

```
switch# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)#
```

Configuration Mode Commands and Submodes

Here is a list of configuration mode commands:

```
switch# config t
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# ?
Configure commands:
  aaa                Configure aaa functions
  arp                [no] remove an entry from the ARP cache
  asm                Configure ASM Modules
  banner             Configure banner message
  boot              Configure boot variables
  callhome           Enter the callhome configuration mode
  cdp                CDP Configuration parameters
  cfs                CFS configuration commands
  cimserver          Modify cimserver configuration
  clock              Configure time-of-day clock
  crypto             Set crypto settings
  device-alias       Device-alias configuration commands
  do                 EXEC command
  dpvm              Configure Dynamic Port Vsan Membership
  end                Exit from configure mode
  exit              Exit from configure mode
  fabric-binding     Fabric Binding configuration
  fc-tunnel          Configure fc-tunnel
  fcalias            Fcalias configuration commands
  fcanalyzer         Configure cisco fabric analyzer
  fcc                Configure FC Congestion Control
  fcdomain           Enter the fcdomain configuration mode
  fcdroplateness     Configure switch or network latency
  fcflow             Configure fcflow
  fcid-allocation    Add/remove company id(or OUIs) from auto area list
  fcinterop          Interop commands
  fcip              Enable/Disable FCIP
  fcns              Name server configuration
  fcroute            Configure FC routes
  fcrxbcredit        Enable extended rx b2b credit configuration
  fcs                Configure Fabric Config Server
  fcsp              Config commands for FC-SP
  fctimer            Configure fibre channel timers
  fdmi              Config commands for FDMI
  ficon              Configure ficon information
  fspf              Configure fspf
  in-order-guarantee Set in-order delivery guarantee
  interface          Select an interface to configure
  ip                 Configure IP features
  iscsi              Enable/Disable iSCSI
```


Send documentation comments to mdsfeedback-doc@cisco.com.

| | |
|----------------------|---|
| ivr | Config commands for IVR |
| kernel | Kernel options |
| line | Configure a terminal line |
| logging | Modify message logging facilities |
| mcast | Configure multicast |
| nasb | Configure Third-Party Copy Functionality |
| no | Negate a command or set its defaults |
| ntp | NTP Configuration |
| port-security | Configure Port Security |
| port-track | Configure Switch port track config |
| power | Configure power supply |
| poweroff | Poweroff a module in the switch |
| qos | QoS Configuration commands |
| radius | Configure RADIUS configuration distribution |
| radius-server | Configure RADIUS related parameters |
| rib | Configure RIB parameters |
| rmon | Remote Monitoring |
| role | Configure roles |
| rscn | Config commands for RSCN |
| san-ext-tuner | Enable/Disable San Extension Tuner tool |
| santap | Enter SanTap configuration |
| scheduler | Config commands for scheduler |
| scsi-flow | SCSI Flow configuration |
| snmp-server | Configure snmp server |
| span | Enter SPAN configuration mode |
| ssh | Configure SSH parameters |
| ssm | Config commands for SSM (Storage Services Module) |
| switchname | Configure system's network name |
| system | System config command |
| tacacs+ | Enable tacacs+ |
| telnet | Enable telnet |
| tlport | Configure TL Port information |
| trunk | Configure Switch wide trunk protocol |
| username | Configure user information. |
| vsan | Enter the vsan configuration mode |
| wwn | Set secondary base MAC addr and range for additional WWNs |
| zone | Zone configuration commands |
| zone-attribute-group | Zone attribute group commands |
| zoneset | Zoneset configuration commands |

Configuration mode, also known as terminal configuration mode, has several submodes. Each of these submodes places you deeper in the prompt hierarchy. When you type **exit**, the switch backs out one level and returns you to the previous level. When you type **end**, the switch backs out to the user EXEC level. You can also type **Ctrl-Z** in configuration mode as an alternative to typing **end**.



Note

In configuration mode, you can alternatively enter

- **Ctrl-Z** instead of the **end** command, and
- **Ctrl-G** instead of the **exit** command

You can execute an EXEC mode command from a configuration mode or submode prompt. You can issue this command from any submode within the configuration mode. When in configuration mode (or in any submode), enter the **do** command along with the required EXEC mode command. The entered command is executed at the EXEC level and the prompt resumes its current mode level.

```
switch(config)# do terminal session-timeout 0
switch(config)#
```

In this example, **terminal session-timeout** is an EXEC mode command—you are issuing an EXEC mode command using the configuration mode **do** command.

Send documentation comments to mdsfeedback-doc@cisco.com.

The **do** command applies to all EXEC mode commands other than the **end** and **exit** commands. You can also use the help (?) and command completion (**Tab**) features for EXEC commands when issuing a **do** command along with the EXEC command.

Table 2-2 lists some useful command keys that can be used in both EXEC and configuration modes:

Table 2-2 Useful Command Key Description

| Command | Description |
|-----------------|--|
| Ctrl-P | Up history. |
| Ctrl-N | Down history. |
| Ctrl-R | Refreshes the current line and reprints it. |
| Ctrl-X-H | List history. |
| Alt-P | History search backwards. Note The difference between Tab completion and Alt-P or Alt-N is that Tab completes the current word while Alt-P and Alt-N completes a previously entered command. |
| Alt-N | History search forwards. |
| Ctrl-G | Exit. |
| Ctrl-Z | End. |
| Ctrl-L | Clear screen. |

CLI Command Navigation

To redisplay a command you previously entered, press the **Up Arrow** key. You can continue to press the **Up Arrow** key to see more previously issued commands. Similarly, you can press the **Down Arrow**, **Right Arrow**, **Left Arrow**, and **Delete** keys to navigate through the command history and to modify an existing command string.

Getting Help

In any command mode, you can get a list of available commands by entering a question mark (?).

```
switch# ?
```

To obtain a list of commands that begin with a particular character sequence, type in those characters followed immediately by the question mark (?). Do not include a space.

```
switch# co?
configure copy
```

To list keywords or arguments, enter a question mark in place of a keyword or argument. Include a space before the question mark. This form of help is called command syntax help, because it reminds you which keywords or arguments are applicable based on the commands, keywords, and arguments you have already entered.

```
switch# config ?
terminal Configure the system from the terminal
```

Send documentation comments to mdsfeedback-doc@cisco.com.



Tip

If you are having trouble entering a command, check the system prompt and enter the question mark (?) for a list of available commands. You might be in the wrong command mode or using incorrect syntax.

Command Completion

In any command mode, you can begin a particular command sequence and immediately press the **Tab** key to complete the rest of the command.

```
switch(config)# ro <Tab>
switch(config)# role <Tab>
switch(config)# role name
```

This form of help is called command completion, because it completes a word for you. If several options are available for the typed letters, all options that match those letters are presented:

```
switch(config)# fc <Tab>
fcalias          fcdomain          fcs
fcalyzer         fcdroplacency     fcns          fctimer
fcc              fcinterop         fcroute
switch(config)# fcd <Tab>
fcdomain         fcdroplacency
switch(config)# fcd <Tab>
switch(config)# fcdomain
```

File System Completion

As of SAN-OS Release 2.0(1b), you can use the **Tab** key to complete schemes, servers, and file names available in the file system.

For example,

```
switch# cd bootflash:<Tab>
bootflash:          bootflash://sup-1/          bootflash://sup-remote/
bootflash:///        bootflash://sup-2/          bootflash://sup-standby/
bootflash://module-5/ bootflash://sup-active/
bootflash://module-6/ bootflash://sup-local/

switch# bootflash://mo<Tab>
switch# bootflash://module-6/
```

The no and Default Forms of Commands

You can issue the **no** form of any command to perform the following actions:

- Undo a wrongly issued command.

If you issue the **zone member** command, you can undo the results:

```
switch(config)# zone name test vsan 1
switch(config-zone)# member pwn 12:12:12:12:12:12
switch(config-zone)# no member pwn 12:12:12:12:12:12
WARNING: Zone is empty. Deleting zone test. Exit the submode.
switch(config-zone)#
```

Send documentation comments to mdsfeedback-doc@cisco.com.

- Delete a created facility.

If you want to delete a zone that you created:

```
switch(config)# zone name test vsan 1
switch(config-zone)# exit
switch(config)# no zone name test vsan 1
switch(config)#
```

You cannot delete a zone facility called test while residing in it. You must first exit the zone submode and return to configuration mode.

CLI Command Configuration Options

You can configure the software in one of two ways:

- You can create the configuration for the switch interactively by issuing commands at the CLI prompt.
- You can create an ASCII file containing a switch configuration and then load this file on the required system. You can then use the CLI to edit and activate the file (see the [“Working with Configuration Files”](#) section on page 4-24).

Displaying the Switch Configuration

You can view the ASCII form of the configuration file when required. To view the current configuration tree from the EXEC prompt, issue the **show running-config** command. If the running configuration is different from the startup configuration, issue the **show startup-config** command to view the ASCII version of the current startup configuration that was used to boot the switch if a **copy run start** command was not issued after the reboot. Use the **show startup** command to view the contents of the current startup configuration.

You can also gather specific information on the entire switch configuration by issuing the relevant **show** commands. Configurations are displayed based on a specified feature, interface, module, or VSAN. Available **show** commands for each feature are briefly described in this section and listed at the end of each chapter.

Examples 2-2 to 2-8 display a few **show** command examples.

Example 2-2 *Displays Details on the Specified Interface*

```
switch# show interface fc1/1
fc1/1 is up
  Hardware is Fibre Channel, 20:01:ac:16:5e:4a:00:00
  vsan is 1
  Port mode is E
  Speed is 1 Gbps
  Beacon is turned off
  FCID is 0x0b0100
    0 frames input, 0 bytes, 0 discards
    0 runts, 0 jabber, 0 too long, 0 too short
    0 input errors, 0 CRC, 0 invalid transmission words
    0 address id, 0 delimiter
    0 EOF abort, 0 fragmented, 0 unknown class
    0 frames output, 0 bytes, 0 discards
  Received 0 OLS, 0 LRR, 0 NOS, 0 loop inits
  Transmitted 0 OLS, 0 LRR, 0 NOS, 0 loop inits
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 2-3 Displays the Software and Hardware Version

```
switch# show version
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2003, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
Cisco Systems, Inc. and/or other third parties and are used and
distributed under license. Some parts of this software are covered
under the GNU Public License. A copy of the license is available
at http://www.gnu.org/licenses/gpl.html.

Software
  BIOS:          version 1.0.8
  loader:        version 1.1(2)
  kickstart:     version 2.0(1) [build 2.0(0.6)] [gdb]
  system:        version 2.0(1) [build 2.0(0.6)] [gdb]

  BIOS compile time:      08/07/03
  kickstart image file is: bootflash://m9500-sflek9-kickstart-mzg.2.0.0.6.bin
  kickstart compile time: 10/25/2010 12:00:00
  system image file is:   bootflash://m9500-sflek9-mzg.2.0.0.6.bin
  system compile time:    10/25/2020 12:00:00

Hardware
  RAM 1024584 kB

  bootflash: 1000944 blocks (block size 512b)
  slot0:      0 blocks (block size 512b)

172.22.92.181 uptime is 0 days 2 hours 18 minute(s) 1 second(s)

Last reset at 970069 usecs after Tue Sep 16 22:31:25 1980
  Reason: Reset Requested by CLI command reload
  System version: 2.0(0.6)
  Service:
```

Example 2-4 Displays the Running Configuration

```
switch# show running
Building Configuration ...
  interface fc1/1
  interface fc1/2
  interface fc1/3
  interface fc1/4
  interface mgmt0
ip address 172.22.95.112 255.255.255.0
no shutdown
vsan database
boot system bootflash:system-237; sup-1
boot kickstart bootflash:boot-237 sup-1
callhome
ip default-gateway 172.22.95.1
switchname switch
trunk protocol enable
username admin password 5 /AFDAMD4B2xK2 role network-admin
```

Example 2-5 Displays the Difference Between the Running and Startup Configuration

```
switch# show running diff
Building Configuration ...
*** Startup-config
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```

--- Running-config
***** 1,16 ****
    fcip enable
    ip default-gateway 172.22.91.1
    iscsi authentication none
    iscsi enable
! iscsi import target fc
    iscsi virtual-target name vt
        pWWN 21:00:00:04:cf:4c:52:c1
    all-initiator-permit
--- 1,20 ----
    fcip enable
+ aaa accounting logsize 500
+
+
+
    ip default-gateway 172.22.91.1
    iscsi authentication none
    iscsi enable
! iscsi initiator name junk
    iscsi virtual-target name vt
        pWWN 21:00:00:04:cf:4c:52:c1
    all-initiator-permit

```

Example 2-6 *Displays the Configuration for a Specified Interface*

```

switch# show running interface fc2/9
interface fc2/9
switchport mode E
no shutdown

```



Note

The **show running interface** command is different from the **show interface** command.

Example 2-7 *Displays the Configuration for all Interfaces in a 16-Port Module*

```

switch# show running interface fc2/10 - 12
interface fc2/10
switchport mode E
no shutdown

interface fc2/11
switchport mode E
no shutdown

interface fc2/12
switchport mode FL
no shutdown

```

Example 2-8 *Displays the Configuration Per VSAN*

```

switch# show running vsan 1
Building Configuration ...
zone name m vsan 1
    member pwwn 21:00:00:20:37:60:42:5c
    member pwwn 21:00:00:20:37:4b:00:a2
zoneset name m vsan 1
    member m
zoneset activate name m vsan 1

```

Send documentation comments to mdsfeedback-doc@cisco.com.

Saving a Configuration

Use the **copy running-config startup-config** command to save the new configuration into nonvolatile storage. Once this command is issued, the running and the startup copies of the configuration are identical.

See the “Copying Files” section on page 4-28 and the “Preserving Module Configuration” section on page 7-7.

Clearing a Configuration

Use the **write erase** command to clear a startup configuration. Once this command is issued, the switch's startup configuration reverts to factory defaults. The running configuration is not affected.



Caution

The **write erase** command erases the entire startup configuration with the exception of any configuration that affects the loader functionality.

The **write erase boot** command only erases the configuration that affects the loader functionality. The loader functionality configuration includes the boot variables and the mgmt0 IP configuration information (IP address, netmask, and default gateway).

```
switch# write erase boot
```

This command will erase the boot variables and the ip configuration of interface mgmt 0

Displaying Users

Use the **show users** command to display all users currently accessing the switch.

```
switch# show users
admin pts/7 Jan 12 20:56 (10.77.202.149)
admin pts/9 Jan 12 23:29 (modena.cisco.com)
admin pts/11 Jan 13 01:53 (dhcp-171-71-49-49.cisco.com)
```

Sending Messages to Users

Use the **send** command to send a message to all active CLI users currently using the switch. This message is restricted to 80 alphanumeric characters with spaces.

This example sends a warning message to all active users about the switch being shut down.

```
switch# send Shutting down the system in 2 minutes. Please log off.
```

```
Broadcast Message from admin@excal-112
(/dev/pts/3) at 16:50 ...
Shutting down the system in 2 minutes. Please log off.
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Using the ping Command

Use the **ping** command to verify the connectivity of a remote host or server by sending echo messages.

The syntax for this command is **ping** <host or ip address>.

```
switch# ping 198.133.219.25
PING 198.133.219.25 (198.133.219.25) 56(84) bytes of data.
64 bytes from 198.133.219.25: icmp_seq=1 ttl=245 time=0.856 ms
64 bytes from 198.133.219.25: icmp_seq=2 ttl=245 time=1.02 ms

--- 198.133.219.25 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 999ms
rtt min/avg/max/mdev = 0.856/0.941/1.027/0.090 ms
```

To abnormally terminate a ping session, type the **Ctrl-C** escape sequence.

Using the Extended ping Command

As of SAN-OS Release 2.0(1b), the **ping** command provides additional options to verify the connectivity of a remote host or server. To specify these additional parameters, just type **ping** at the CLI switch prompt and press enter.

[Table 2-3](#) summarizes the syntax and the defaults.

Table 2-3 Options and Defaults for the ping Command

| Option | Description | Default |
|-----------------------------|---|----------------|
| Target IP address | The IP address or host name of the destination node to ping. | Not applicable |
| Repeat count | The number of ping packets to be sent to the destination address. | 5 packets |
| Datagram size | The size of each ping packet in bytes. | 100 bytes |
| Timeout in seconds | The timeout interval before the ping command is terminated. | 2 seconds |
| Extended commands | Specifies if a series of additional commands appear. | No |
| Sweep range of sizes | The sizes of the echo packets being sent. This option determines the minimum sizes of the MTUs configured on the nodes along the path to the destination address. You can then reduce packet fragmentation performance problems. (see the “Configuring the MTU Frame Size” section on page 28-6). | No |
| Source address or interface | The numeric IP address or the name of the source interface. | Not applicable |
| Type of service | The Quality of Service (QoS) in Internet Control Message Protocol (ICMP) datagrams (see the “QoS” section on page 32-3). | 0 |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 2-3 Options and Defaults for the ping Command (continued)

| Option | Description | Default |
|-------------------------|---|---------|
| Set DF bit in IP header | The Path MTU Discovery strategy (see the “Configuring the MTU Frame Size” section on page 28-6). | No |
| Data pattern | You may specify up to 16 bytes to pad the outgoing packet. This padding is useful when diagnosing data-dependent problems in a network. For example, <code>ff</code> fills the outgoing packet with all ones. | 0xABCD |

The syntax for this command is as follows:

```
switch# ping
Target IP address: 198.133.219.25
Target IP address: 198.133.219.25
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface:
Type of service [0]:
Set DF bit in IP header [n]:
Data pattern [0xABCD]:
Sweep range of sizes [n]:
PATTERN: 0xabcd
PING 198.133.219.25 (198.133.219.25) 100(128) bytes of data.
108 bytes from 198.133.219.25: icmp_seq=1 ttl=245 time=0.600 ms
108 bytes from 198.133.219.25: icmp_seq=2 ttl=245 time=0.614 ms
108 bytes from 198.133.219.25: icmp_seq=3 ttl=245 time=0.872 ms
108 bytes from 198.133.219.25: icmp_seq=4 ttl=245 time=0.558 ms
108 bytes from 198.133.219.25: icmp_seq=5 ttl=245 time=0.570 ms

--- 198.133.219.25 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 7996ms
rtt min/avg/max/mdev = 0.558/0.642/0.872/0.120 ms
```

To abnormally terminate a ping session, type the **Ctrl-C** escape sequence.

Send documentation comments to mdsfeedback-doc@cisco.com.

Using traceroute

Use the **traceroute** command to prints the routes taken by a specified host or IP address.

The syntax for this command is **traceroute** *<host or ip address>*.

```
switch# traceroute www.cisco.com
Tracing route to www.cisco.com [198.133.219.25] 30 hops max, 38 byte packets
 1  bras3-10.pltnca.sbcglobal.net [151.164.184.79] 30 ms 30 ms 20 ms
 2  dist2-vlan50.pltn13.pbi.net [64.164.97.67] 20 ms 20 ms 30 ms
 3  bb2-g1-1.pltn13.pbi.net [67.116.251.194] 20 ms 20 ms 20 ms
 4  bb1-p12-0.pltn13.pbi.net [151.164.40.17] 20 ms 21 ms 20 ms
 5  bb2-p13-0.sntc01.pbi.net [151.164.191.65] 20 ms 20 ms 30 ms
 6  ex1-p3-0.eqsjca.sbcglobal.net [64.161.1.54] 20 ms 20 ms 30 ms
 7  sl-st20-sj-0-0.sprintlink.net [144.223.242.81] 20 ms 20 ms 30 ms
 8  sl-bb25-sj-10-0.sprintlink.net [144.232.20.62] 20 ms 30 ms 20 ms
 9  sl-gw11-sj-10-0.sprintlink.net [144.232.3.134] 70 ms 30 ms 30 ms
10  sl-ciscopsn2-11-0-0.sprintlink.net [144.228.44.14] 20 ms 30 ms 20 ms
11  sjce-dmzbb-gw1.cisco.com [128.107.239.89] 20 ms 30 ms 30 ms
12  sjck-dmzdc-gw1.cisco.com [128.107.224.69] 20 ms 30 ms 20 ms
13  www.cisco.com (198.133.219.25) 2.496 ms * 2.135 ms
```

To abnormally terminate a traceroute session, enter **Ctrl-C**.

Setting the Shell Timeout

Use the **exec-timeout** command in configuration mode to configure the lifetime of all terminal sessions on that switch. When the time limit configured by this command is exceeded, the shell exits and closes that session. The syntax for this command is **exec-timeout** *minutes*.

The default is 30 minutes. You can configure different timeout values for a console or a virtual terminal line (VTY) session. You can set the **exec-timeout** value to 0 to disable this feature so the session remains active until you exit the switch. This change is saved in the configuration file.

- From the console:

```
switch(config)# line console
switch(config-console)# exec-timeout 60
```

Specifies the current console shell timeout to be 60 minutes.

- From a VTY session (Telnet or SSH):

```
switch(config)# line vty
switch(config-line)# exec-timeout 60
```

Specifies the current console shell timeout to be 60 minutes.

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying VTY Sessions

Use the **show line** command to display all configured VTY sessions:

```
switch# show line
line Console:
  Speed:          9600 bauds
  Databits:       8 bits per byte
  Stopbits:       1 bit(s)
  Parity:         none
  Modem In: Disable
  Modem Init-String -
    default : ATE0Q1&D2&C1S0=1\015
  Statistics: tx:5558511 rx:5033958 Register Bits:RTS|CTS|DTR|DSR|CD|RI
line Aux:
  Speed:          9600 bauds
  Databits:       8 bits per byte
  Stopbits:       1 bit(s)
  Parity:         none
  Modem In: Disable
  Modem Init-String -
    default : ATE0Q1&D2&C1S0=1\015
  Hardware Flowcontrol: ON
  Statistics: tx:35 rx:0 Register Bits:RTS|DTR
```

Clearing VTY Sessions

Use the **clear line** command to close a specified VTY session:

```
switch# clear line Aux
```

Setting the Terminal Timeout

Use the **terminal session-timeout** command in EXEC mode to configure the automatic logout time for the current terminal session on that switch. When the time limit configured by this command is exceeded, the switch closes that session and exits.

The syntax for this command from is **terminal session-timeout** *minutes*.

The default is 30 minutes. You can set the **terminal session-timeout** value to 0 to disable this feature so the terminal remains active until you choose to exit the switch. This change is not saved in the configuration file.

```
switch# terminal session-timeout 600
```

Specifies the terminal timeout to be 600 minutes for the current session.

Send documentation comments to mdsfeedback-doc@cisco.com.

Setting the Terminal Type

Use the **terminal terminal-type** command in EXEC mode to specify the terminal type for a switch:

The syntax for this command is **terminal terminal-type** *terminal-type*.

```
switch# terminal terminal-type vt100
```

Specifies the terminal type. The *terminal-type* string is restricted to 80 characters and must be a valid type (for example vt100 or xterm). If a Telnet or SSH session specifies an unknown terminal type, the switch uses the vt100 terminal by default.

Setting the Terminal Length

Use the **terminal length** command in EXEC mode to set the terminal screen length for the current session. This command is specific to only the console port. Telnet and SSH sessions set the length automatically.

The syntax for this command is **terminal length** *lines*.

```
switch# terminal length 20
```

Sets the screen length for the current session to 20 lines for the current terminal session. The default is 24 lines.

Setting the Terminal Width

Use the **terminal width** command in EXEC mode to set the terminal screen width for the current session. This command is specific to only the console port. Telnet and SSH sessions set the width automatically.

The syntax for this command is **terminal width** *columns*.

```
switch# terminal width 86
```

Sets the screen length for the current session to 86 columns for the current terminal session. The default is 80 columns.

Displaying Terminal Settings

Use the **show terminal** command to display the terminal settings for the current session:

```
switch# show terminal
TTY: Type: "vt100"
Length: 24 lines, Width: 80 columns
Session Timeout: 525600 minutes
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Configuring the Switch Banner Message

As of Cisco MDS SAN-OS Release 1.3(4), you can issue the **banner motd** command in configuration mode to configure a message of the day (MOTD).

The syntax for this command from is **banner motd** [*delimiting-character message delimiting-character*]

The following example configures a banner message with the following text “Testing the MOTD Feature.”

```
switch# config t
switch(config)# banner motd # Testing the MOTD Feature. #
```

The message is restricted to 40 lines with a maximum of 80 characters in each line.

Use the **show banner motd** command to display the configured banner message:

The following example displays the configured banner message.

```
switch# show banner motd
Testing the MOTD Feature
```

The configured MOTD banner is displayed before the login prompt on the terminal whenever a user logs in to a Cisco MDS 9000 Family switch.

```
Testing the MOTD Feature
switch login:
```

Follow these guidelines when choosing your delimiting character:

- Do not use the *delimiting-character* in the *message* string.
- Do not use " and % as delimiter s.

You can include tokens in the form \$ (token) in the message text. Tokens will be replaced with the corresponding configuration variable. For example:

- \$(hostname) displays the host name for the switch
- \$(line) displays the vty or tty line or name

The following example spans multiple lines and uses tokens to configure the banner message:

```
switch# config t
switch(config)# banner motd #
Enter TEXT message. End with the character '#'.
Welcome to switch $(hostname).
Your tty line is $(line).
#
```

Send documentation comments to mdsfeedback-doc@cisco.com.

About Flash Devices

Every switch in the Cisco MDS 9000 Family contains one internal bootflash (see [Figure 2-2](#)). The Cisco MDS 9500 Series additionally contains one external CompactFlash called slot0 (see [Figure 2-2](#) and [Figure 2-3](#)).

Figure 2-2 Flash Devices in the Cisco MDS 9000 Supervisor Module





Figure 2-3 External CompactFlash in the Cisco MDS 9000 Supervisor Module



Internal bootflash:

All switches in the Cisco MDS 9000 Family have one internal bootflash: that resides in the supervisor or switching module. You have access to two locations within the internal bootflash: file system.

- The volatile: file system provides temporary storage, and it is also the default location for file system commands. Files in temporary storage (volatile:) are erased when the switch reboots.
- The bootflash: (nonvolatile storage) file system provides permanent storage. The files in bootflash: are preserved through reboots and power outages.

External CompactFlash (Slot0:)

Cisco MDS 9500 Series directors contain an additional external CompactFlash referred to as the slot0: file system.

The external CompactFlash, an optional device for MDS 9500 Series directors, can be used for storing software images, logs, and core dumps.

Send documentation comments to mdsfeedback-doc@cisco.com.

Formatting Flash Devices and File Systems

By formatting a Flash devices or a file system, you are clearing out the contents of the device or the file system and restoring it to its factory-shipped state.

See the “About Flash Devices” section on page 2-20 and the “Using the File System” section on page 2-22.

Initializing Internal bootflash:

When a switch is shipped, the **init system** command is already performed and you do not need to issue it again. Initializing the switch resets the entire internal Flash device and erases all data in the bootflash: file system. The internal Flash device is composed of several file systems with bootflash: being one of them. All files in bootflash: are erased and you must download the system and kickstart images again. After issuing an **init system** command, you do not have to format the bootflash: again because bootflash: is automatically formatted.



Note

The **init system** command also installs a new loader from the existing (running) kickstart image. You can access this command from the `switch(boot)#` prompt (see Chapter 6, “Software Images”).

If bootflash: is found corrupted during a boot sequence, you will see the following message:

```
ERROR:bootflash: has unrecoverable error; please do "format bootflash:"
```

Use the **format bootflash:** command to only format the bootflash: file system. You can issue the **format bootflash:** command from either the `switch#` or the `switch(boot)#` prompts.

If you issue the **format bootflash:** command, you must download the kickstart and system images again.

Formatting External CompactFlash

Be sure to format an external CompactFlash device before using it to save files or images.

You can verify if the external CompactFlash device is formatted by inserting it into slot0: and issuing the **dir slot0:** command.

- If the external CompactFlash device is already formatted, you can see file system usage information (along with any existing files).
- If the external CompactFlash device is unformatted (corrupted), you will see the following message:

```
Device unavailable
```

In this case, you need to format the CompactFlash device using the **format slot0:** command.



Note

The slot0: file system cannot be accessed from the standby `loader>` prompt or the `switch(boot)#` prompt, if the disk is inserted after booting the switch.

Send documentation comments to mdsfeedback-doc@cisco.com.



Caution

The Cisco SAN-OS software only supports CompactFlash devices that are certified by Cisco Systems and formatted using Cisco MDS switches. Using uncertified CompactFlash devices may result in unpredictable consequences; formatting CompactFlash devices using other platforms may result in errors.

Using the File System

The switch provides the following useful functions to help you manage software image files and configuration files:

- [Setting the Current Directory, page 2-22](#)
- [Displaying the Current Directory, page 2-23](#)
- [Listing the Files in a Directory, page 2-23](#)
- [Creating a Directory, page 2-23](#)
- [Deleting an Existing Directory, page 2-24](#)
- [Moving Files, page 2-24](#)
- [Copying Files, page 2-25](#)
- [Deleting Files, page 2-25](#)
- [Displaying File Contents, page 2-25](#)
- [Saving Command Output to a File, page 2-26](#)
- [Compressing and Uncompressing Files, page 2-26](#)
- [Displaying the Last Lines in a File, page 2-27](#)
- [Executing Commands Specified in a Script, page 2-27](#)
- [Setting the Delay Time, page 2-28](#)

Setting the Current Directory

The **cd** command changes the current directory level to a specified directory level. CLI defaults to the volatile: file system. This command expects a directory name input.



Tip

Any file saved in the volatile: file system is erased when the switch reboots.

The syntax for this command is **cd** *directory name*

This example changes the current directory to the mystorage directory that resides in the slot0 file system:

```
switch# cd slot0:mystorage
```

This example changes the current directory to the mystorage directory that resides in the current directory.

```
switch# cd mystorage
```


Send documentation comments to mdsfeedback-doc@cisco.com.

If the current directory is `slot0:mydir`, this command changes the current directory to `slot0:mydir/mystorage`.

Displaying the Current Directory

The **pwd** command displays the current directory location. This example changes the directory and displays the current directory.

```
switch# cd bootflash:
switch# pwd
bootflash:
```



Note

If you issue this command from the active supervisor module in a Cisco MDS 9500 Series (for example, `module-5`), then you cannot change the current working directory to the bootflash: of `module-6`. See the “Supervisor Modules” section on page 7-2.

Displaying File Checksums

The **show file file md5sum** command provides the MD5 checksum of file. MD5 is an electronic fingerprint for the file. MD5 is the latest implementation of the internet standards described in RFC 1321 and is useful for data security as well as integrity.

The **show file file cksum** command provides the checksum of file. The checksum values compute a cyclic redundancy check (CRC) for each named file. Use this command to verify that are not corrupted—compare the checksum output for the received file against the checksum output for the original file.

This example provides the output of the **show file** command when a file is specified.

```
switch# show file bootflash://sup-1/ultimate_file.tar cksum
2569913991

switch# show file bootflash://sup-1/ultimate_file.tar md5sum
52479aae2dce1fd849b6f4916d750392
```

Listing the Files in a Directory

The **dir** command displays the contents of the current directory or the specified directory. The syntax for this command is **dir** *directory or file name*.

This example shows how to list the files on the default volatile: file system:

```
switch# dir
Usage for volatile: filesystem
          0 bytes total used
    20971520 bytes free
    20971520 bytes available
```

Creating a Directory

The **mkdir** command creates a directory at the current directory level or at a specified directory level.

The syntax for this command is **mkdir** *directory name*.

Send documentation comments to mdsfeedback-doc@cisco.com.

This example creates a directory called test in the slot0 directory.

```
switch# mkdir slot0:test
```

This example creates a directory called test at the current directory level.

```
switch# mkdir test
```

If the current directory is slot0:mydir, this command creates a directory called slot0:mydir/test.

Deleting an Existing Directory

The **rmdir** command deletes an existing directory at the current directory level or at a specified directory level. The directory must be empty to be deleted.

The syntax for this command is **rmdir** *directory name*.

This example deletes the directory called test in the slot0 directory.

```
switch# rmdir slot0:test
This is a directory. Do you want to continue (y/n)? [y] y
```

As of Cisco SAN-OS Release 2.0(1b), the **delete** command is also capable of deleting empty and non-empty directories. When you issue this command a warning is displayed to confirm your intent to delete the directory

This example deletes the directory called test at the current directory level.

```
switch# rmdir test
This is a directory. Do you want to continue (y/n)? [y] y
```

If the current directory is slot0:mydir, this command deletes the slot0:mydir/test directory.

Moving Files

The **move** command removes a file from the source directory and places it in the destination directory. If a file with the same name already exists in the destination directory, that file is overwritten by the moved file.

This example moves the file called samplefile from the root directory of the slot0: file system to the mystorage directory.

```
switch# move slot0:samplefile slot0:mystorage/samplefile
```

This example moves a file from the current directory level.

```
switch# move samplefile mystorage/samplefile
```

If the current directory is slot0:mydir, this command moves slot0:mydir/samplefile to slot0:mydir/mystorage/samplefile.

Send documentation comments to mdsfeedback-doc@cisco.com.

Copying Files

The **copy** command copies a file.

This example copies the file called samplefile from the root directory of the slot0: file system to the mystorage directory.

```
switch# copy slot0:samplefile slot0:mystorage/samplefile
```

This example copies a file from the current directory level.

```
switch# copy samplefile mystorage/samplefile
```

If the current directory is slot0:mydir, this command copies slot0:mydir/samplefile to slot0:mydir/mystorage/samplefile.

You can also use the **copy** command to upload and download files from the slot0: or bootflash: file system to or from a FTP, TFTP, SFTP, or SCP server (see the [“Copying Files” section on page 4-28](#)).

Deleting Files

The **delete** command deletes a specified file or the specified directory and all its contents (see the [“Deleting Files” section on page 4-33](#)).

This example shows how to delete a file from the current working directory:

```
switch# delete dns_config.cfg
```

This example shows how to delete a file from an external CompactFlash (slot0):

```
switch# delete slot0:dns_config.cfg
```

This example deletes the entire my-dir directory and all its contents:

```
switch# delete bootflash:my-dir
```



Caution

If you specify a directory, the **delete** command deletes the entire directory and all its contents.

Displaying File Contents

The **show file** command displays the contents of a specified file in the file system.

The syntax for this command is **show file** *file_name*

This example displays the contents of the test file that resides in the slot0 directory.

```
switch# show file slot0:test
config t
Int fc1/1
no shut
end
show int
```

This example displays the contents of a file residing in the current directory.

```
switch# show file myfile
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Saving Command Output to a File

You can force all screen output to go to a file by appending `> filename` to any command. For example, enter **show interface > samplefile** at the EXEC mode switch prompt to save the interface configuration to *samplefile*—a file created at the same directory level. At the EXEC mode switch prompt, issue a **dir** command to view all files in this directory, including the recently saved *samplefile*. See [Chapter 4, “Initial Configuration,”](#) for information on saving and copying configuration files, and [Chapter 6, “Software Images,”](#) for information on saving and copying software images.



Note

Redirection is allowed only if the current directory is on the `volatile:` (default) or `slot0:` file systems. Redirection is not allowed if the current directory is on the `bootflash:` file system. The current directory can be viewed using the **pwd** command and changed using the **cd** command.

Compressing and Uncompressing Files

The **gzip** command compresses (zips) the specified file using LZ77 coding.

This example directs the output of the `show tech-support` command to a file (*Samplefile*) and then zips the file and displays the difference in the space used in the `volatile:` directory:

```
switch# show tech-support > Samplefile
Building Configuration ...
switch# dir
    1525859      Jul 04 00:51:03 2003 Samplefile
Usage for volatile://
    1527808 bytes used
    19443712 bytes free
    20971520 bytes total
switch# gzip volatile:Samplefile
switch# dir
    266069      Jul 04 00:51:03 2003 Samplefile.gz
Usage for volatile://
    266240 bytes used
    20705280 bytes free
    20971520 bytes total
```

The **gunzip** command uncompresses (unzips) LZ77 coded files.

This example unzips the file that was compressed in the previous example:

```
switch# gunzip samplefile
switch# dir
    1525859      Jul 04 00:51:03 2003 Samplefile
Usage for volatile://
    1527808 bytes used
    19443712 bytes free
    20971520 bytes total
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying the Last Lines in a File

The **tail** command displays the last lines (tail end) of a specified file.

The syntax for this command is **tail** *<file name>* [*<number of lines>*]

```
switch# tail mylog 10
```

You see the last 10 lines of the mylog file.

Executing Commands Specified in a Script

The **run-script** command executes the commands specified in a file. To use this command, be sure to create the file and specify commands in the required order.



Note

You cannot create the script files at the switch prompt. You can create the script file on an external machine and copy it the bootflash: directory. This section assumes that the script file resides in the bootflash: directory.

The syntax for this command is **run-script** *file_name*

This example displays the CLI commands specified in the testfile that resides in the slot0 directory.

```
switch# show file slot0:testfile
conf t
interface fc 1/1
no shutdown
end
sh interface fc1/1
```

This file output is in response to the **run-script** command executing the contents in the testfile file:

```
switch# run-script slot0:testfile
'conf t'
Enter configuration commands, one per line. End with CNTL/Z.
'interface fc1/1'
'no shutdown'
'end'
'sh interface fc1/1'
fc1/1 is down (Fcot not present)
  Hardware is Fibre Channel
  Port WWN is 20:01:00:05:30:00:48:9e
  Admin port mode is auto, trunk mode is on
  vsan is 1
  Beacon is turned off
  Counter Values (current):
    0 frames input, 0 bytes, 0 discards
    0 runts, 0 jabber, 0 too long, 0 too short
    0 input errors, 0 CRC, 0 invalid transmission words
    0 address id, 0 delimiter
    0 EOF abort, 0 fragmented, 0 unknown class
    0 frames output, 0 bytes, 0 discards
    Received 0 OLS, 0 LRR, 0 NOS, 0 loop inits
    Transmitted 0 OLS, 0 LRR, 0 NOS, 0 loop inits
  Counter Values (5 minute averages):
...
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Setting the Delay Time

The **sleep** command delays an action by a specified number of seconds.

The syntax for this command is **sleep <seconds>**

```
switch# sleep 30
```

You will see the switch prompt return after 30 seconds.

This command is useful within scripts. For example, if you create a script called test-script:

```
switch# show file slot0:test-script  
discover scsi-target remote  
sleep 10  
show scsi-target disk  
switch# run-script slot0:test-script
```

When you execute the slot0:test-script, the switch software executes the **discover scsi-target remote** command, and then waits for 10 seconds before executing the **show scsi-target disk** command.



Obtaining and Installing Licenses

Licenses are available in all switches in the Cisco MDS 9000 Family. Licensing allows you to access specified premium features on the switch after you install the appropriate license for that feature. Licenses are sold, supported, and enforced as of Cisco MDS SAN-OS Release 1.3(1).

This chapter contains information related to licensing types, options, procedures, installation, and management for the Cisco MDS SAN-OS software.

This chapter includes the following sections:

- [Licensing Terminology, page 3-2](#)
- [Licensing Model, page 3-3](#)
- [Licensing High Availability, page 3-5](#)
- [Options to Install a License, page 3-6](#)
- [Obtaining a Factory-Installed License, page 3-6](#)
- [Performing a Manual Installation, page 3-6](#)
- [Obtaining the License Key File, page 3-7](#)
- [Installing the License Key File, page 3-8](#)
- [Backing Up License Files, page 3-9](#)
- [Identifying License Features in Use, page 3-9](#)
- [Uninstalling Licenses, page 3-9](#)
- [Updating Licenses, page 3-10](#)
- [License Expiry Alerts, page 3-11](#)
- [Grace Period Countdown, page 3-12](#)
- [License Transfers Between Switches, page 3-12](#)
- [Displaying License Information, page 3-12](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

Licensing Terminology

The following terms are used in this chapter:

- **Licensed feature**—Permission to use a particular feature through a license file, a hardware object, or a legal contract. This permission is limited to the number of users, number of instances, time span, and the implemented switch.
- **License expiry**—The time span during which a licensed feature is valid. The software tracks all licenses and sends periodic alerts before shutting down the licensed feature.
- **Counted license**—The number of usage instances for a licensed feature.
- **Licensed application**—A software feature that requires a license to be used.
- **License enforcement**—A mechanism that prevents a feature from being used without first obtaining a license.
- **Node-locked license**—A license that can only be used on a particular switch using the switch's unique host ID.
- **Host IDs**—A unique chassis serial number that is specific to each Cisco MDS switch.
- **Proof of purchase**—A document entitling its rightful owner to use licensed feature(s) on one Cisco MDS switch as described in that document. Also known as the claim certificate.
- **Product Authorization Key (PAK)**—The PAK allows you to obtain a license key from one of the sites listed in the proof of purchase document. After registering at the specified website, you will receive your license key file and installation instructions through e-mail.
- **License key file**—A switch-specific unique file that specifies the licensed features. Each file contains digital signatures to prevent tampering and modification. License keys are required to use a licensed feature. License keys are enforced within a specified time span.
 - License keys are required if your switch is running Cisco MDS SAN-OS Release 1.3 or later.
 - License keys are not required to use licensed features in Cisco MDS SAN-OS Release 1.2 or earlier.
- **Counted license**—The number of licenses issued for a single feature (for example, FCIP). You can increase counted licenses (incremental licenses) should a need arise in the future.
- **Incremental license**—An additional licensed feature that was not in the initial license file. License keys are incremental—if you purchase some features now and others later, the license file and the software detect the sum of all features for the specified switch.
- **Evaluation license**—A temporary license. Evaluation licenses are time bound (valid for a specified number of days) and are not tied to a host ID (switch serial number).
- **Permanent license**—A license that is not time bound (does not have an expiry date) is called a permanent license.
- **Grace period**—The amount of time the features in a license package can continue functioning without a license.
- **Support**—If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco Systems, contact Cisco Technical Support at this URL: <http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

Send documentation comments to mdsfeedback-doc@cisco.com.

Licensing Model

The licensing model defined for the Cisco MDS product line has two options:

- Feature-based licensing: features that are applicable to the entire switch. The cost varies based on a per-switch usage. [Table 3-1](#) lists the feature-based license packages.
- Module-based licensing: features that require additional hardware modules. The cost varies based on a per-module usage. An example is the IPS-8 or IPS-4 module using the FCIP feature.



Note

The FCIP license bundled with Cisco MDS 9216i switch enables FCIP on the two fixed IP services ports only. The features enabled on these ports by the bundled license are identical to the features enabled by the FCIP license on the 14/2-port Multiprotocol Services (MPS-14/2) module, such as the SAN Extension over IP Package. If you install a module with IP ports in the empty slot on the Cisco MDS 9216i, a separate FCIP license is required to enable FCIP on the IP ports of the additional line card.

Table 3-1 **Feature-Based Licenses**

| Feature License | Features |
|---|---|
| Standard package (free—no license required) | <ul style="list-style-type: none"> • FCP, SSH, SFTP, and iSCSI protocols • Fabric manager and remote monitoring (RMON) • VSANs, high availability, PortChannel, and zoning • Fibre Channel Congestion Control (FCC) • Virtual output queuing (VOQ) • Diagnostics (SPAN, RSPAN, and FC Analyzer) • SNMPv3, role-based access control, RADIUS • Call Home and interoperability modes • IP access control lists (ACLs) • Terminal Access Controller Access Control System (TACACS+) • Fabric Device Management Interface (FDMI) • Internet Storage Name Service (iSNS) client. • Cisco Fabric Services (CFS) • Distribute device alias services • Port tracking |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 3-1 Feature-Based Licenses (continued)

| Feature License | Features |
|--|---|
| Enterprise package (ENTERPRISE_PKG) | <ul style="list-style-type: none"> Enhanced security features: <ul style="list-style-type: none"> LUN zoning Read-only zones Port security VSAN-based access control Fibre Channel Security Protocol (FC-SP) authentication IP Security Protocol (IPsec) using the MPS-14/2 module or the Cisco MDS 9216i Switch Advanced traffic engineering—Quality of Service (QoS) Enhanced VSAN routing—inter-VSAN routing (IVR) IVR Network Address Translation (NAT) Zone-based traffic prioritizing Extended credits using the MPS-14/2 module or the Cisco MDS 9216i Switch Zone-based QoS Extended Credits IPsec for FCIP Fibre Channel write acceleration SCSI flow statistics |
| SAN extension over IP package (SAN_EXTN_OVER_IP) | <p>The following features apply to IPS-8, IPS-4, MPS-14/2, and fixed 9216i IP ports</p> <ul style="list-style-type: none"> FCIP FCIP compression FCIP write acceleration FCIP tape acceleration SAN extension tuner features IVR over FCIP IVR NAT Hardware-based FCIP compression FCIP Tape Acceleration SAN Extension Tuner |
| SAN extension over IP package (SAN_EXTN_OVER_IP_IPS4) | |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 3-1 Feature-Based Licenses (continued)

| Feature License | Features |
|--|---|
| Mainframe package (MAINFRAME_PKG) | <ul style="list-style-type: none"> • FICON protocol and CUP management • FICON VSAN and intermixing • Switch cascading • Fabric binding • IBM TotalStorage Virtual Tape Server (VTS) • IBM TotalStorage XRC application |
| Fabric Manager Server package (FM_SERVER_PKG) | <ul style="list-style-type: none"> • Multiple physical fabric management • Centralized fabric discovery services • Continuous MDS health and event monitoring • Long term historical Fibre Channel performance monitoring • Performance reports and charting for hotspot analysis • Web-based operational view • Threshold monitoring • Fabric Manager Web Client for operational view • Performance thresholds |
| Storage Services Enabler package (STORAGE_SERVICES_ENABLER_PKG) | <ul style="list-style-type: none"> • Provides the underlying infrastructure and programmatic interface to enable network-based storage applications when used with the Advanced Services Modules (ASMs) and Storage Services Modules (SSMs). • The network-based storage applications running on the ASM and SSM that require the SSE license are as follows: <ul style="list-style-type: none"> – VERITAS Storage Foundation for Networks • The intelligent fabric applications running on the SSM that require the SSE license are as follows: <ul style="list-style-type: none"> – SANTap – Network-Accelerated Serverless Backup (NASB) – Third-party partner applications |

Licensing High Availability

As with other Cisco MDS SAN-OS features, the licensing feature also maintains the following high availability standards for all switches in the Cisco MDS 9000 Family:

- Installing any license in any switch is a nondisruptive process.
- Installing a license automatically saves a copy of permanent licenses to the chassis in all switches.

Send documentation comments to mdsfeedback-doc@cisco.com.

- Enabling a license feature without a license key starts a counter on the grace period. You then have 60 days (for earlier versions of SAN-OS 1.3(3c) or earlier) or 120 days (for SAN-OS 1.3(4) or later) to install the appropriate license keys or disable the use of that feature. If at the end of the 60 day or 120 day grace period the switch does not have a valid license key for the feature, the feature is automatically disabled by the switch.

Directors in the Cisco MDS 9500 Series have the following additional high availability features:

- The license software runs on both supervisor modules and provides failover protection.
- The license key file is mirrored on both supervisor modules. Even if both supervisor modules fail, the license file continues to function from the version that is available on the chassis.

Options to Install a License

If you have purchased a new switch through either your reseller or through Cisco Systems, you can:

- Obtain a factory-installed license (only applies to new switch orders).
- Perform a manual license installation (applies to existing switches).

Obtaining a Factory-Installed License

You can obtain factory-installed licenses for a new switch.

To obtain a factory-installed license for a new Cisco MDS switch, follow these steps.

Step 1 Contact your reseller or Cisco representative and request this service.



Note If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco Systems, contact Cisco Technical Support at this URL: <http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

Your switch is shipped with the required licenses installed in the system. The proof of purchase document is sent along with the switch.

Step 2 Obtain the host ID from the proof of purchase for future use.

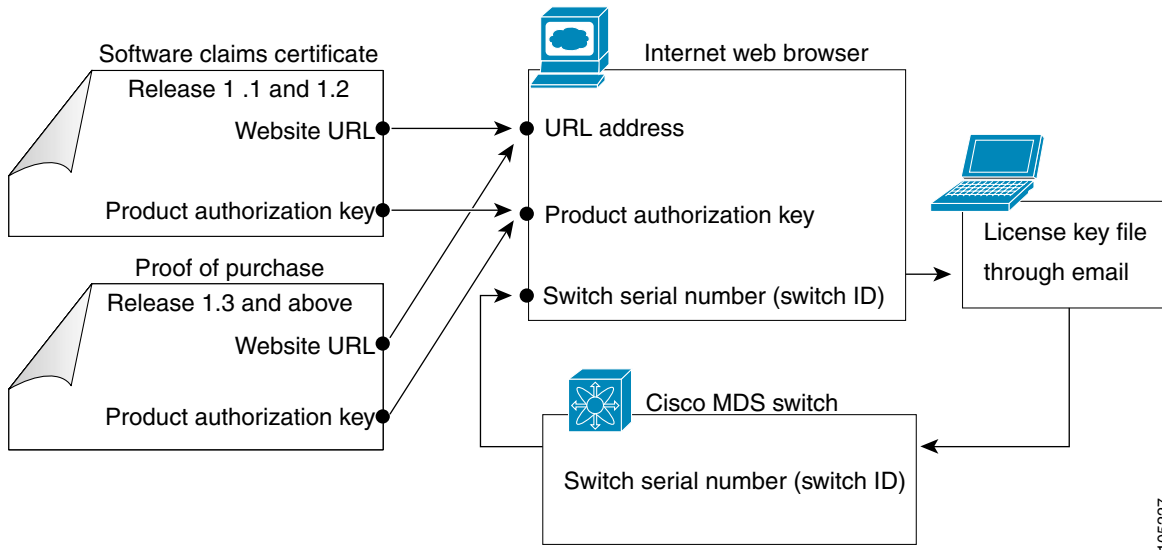
Step 3 Start to use the switch and the licensed features.

Performing a Manual Installation

If you have existing switches or if you wish to install the licenses on your own, you must first obtain the license key file and then install that file in the switch (see [Figure 3-1](#)).

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 3-1 Obtaining a License Key File



105227

Obtaining the License Key File



Note

Refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide* for details on installing automated licenses using the Fabric Manager GUI.

To obtain new or updated license key files using the CLI, follow these steps.

- Step 1** Use the **show license host-id** command to obtain the serial number for your switch. The host ID is also referred to as the switch serial number.

```
switch# show license host-id
License hostid: VDH=FOX064317SQ
```



Tip

Use the entire ID that appears after the colon (:) sign. In this example, the host ID is VDH=FOX064317SQ.

- Step 2** Obtain either your claim certificate or your proof of purchase document. This document accompanies every Cisco MDS switch.
- Step 3** Get the product authorization key (PAK) from either the claim certificate or the proof of purchase document.
- Step 4** Locate the website URL from either the claim certificate or the proof of purchase document.
- Step 5** Access the specified URL that applies to your switch and enter the switch serial number and the PAK. The license key file is sent to you by e-mail. The license key file is digitally signed to only authorize use on the requested switch. The requested features are also enabled once the Cisco SAN-OS software on the specified switch accesses the license key file.

Send documentation comments to mdsfeedback-doc@cisco.com.


Caution

Install the license key file in the specified MDS switch without making any modifications.

A license is either permanent or it expires on a fixed date. If you do not have a license, the grace period for using that feature starts from the first time you start using a feature offered by that license (see the “[License Expiry Alerts](#)” section on page 3-11).

- Step 6** Use the **copy licenses** command in EXEC mode to save your license file to one of two locations—the bootflash: directory or the slot0: device (see the “[Backing Up License Files](#)” section on page 3-9).

Installing the License Key File


Tip

If you need to install multiple licenses in any switch in the Cisco MDS 9000 Family, be sure to provide unique file names for each license key file.

To install a license key file in any switch, follow these steps:

- Step 1** Log into the switch through the console port of the active supervisor.
- Step 2** Perform the installation by issuing the **install license** command on the active supervisor module from the switch console.

```
switch# install license bootflash:license_file.lic
Installing license ..done
```


Note

If you provide a target name for the license key file, the file is installed with the specified name. Otherwise, the file name specified in the license key file is used to install the license.

- Step 3** Exit the switch console and open a new terminal session to view all license files installed on the switch using the **show license** command.

```
switch# show license
Permanent.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT MAINFRAME_PKG cisco 1.0 permanent uncounted \
  HOSTID=VDH=FOX0646S017 \
  NOTICE="<LicFileID></LicFileID><LicLineID>0</LicLineID> \
  <PAK>dummyPak</PAK>" SIGN=EE9F91EA4B64
```


Note

If the license meets all guidelines when the **install license** command is issued, all features and modules continue functioning as configured. This is true for any switch in the Cisco MDS 9000 Family.

Send documentation comments to mdsfeedback-doc@cisco.com.

Backing Up License Files

All installed license files can be backed up as a .tar file in the user specified location. Use the **copy licenses** command in EXEC mode to save your license file to one of three locations—the bootflash: directory or the slot0: device. The following example saves all licenses to a file named Enterprise.tar.

```
switch# copy licenses bootflash:/Enterprise.tar
Backing up license done
```



Tip

We recommend backing up your license files immediately after installing them and just before issuing a **write erase** command.



Caution

If you erase any existing licenses, you can only install them using the **install license** command.

Identifying License Features in Use

When a feature is enabled, it can activate a license grace period. To identify the features active for a specific license using the **show license usage** *license-name* command.

```
switch# show license usage ENTERPRISE_PKG
Application
-----
ivr
qos_manager
-----
```

Uninstalling Licenses

You can only uninstall a permanent license that is not in use. If you try to delete a permanent license that is currently being used, the software rejects the request and issues an error message. Uninstalling an unused license causes the grace period to come into effect. The grace period is counted from the first use of the feature without a license and is reset when a valid license file is installed.



Note

Permanent licenses cannot be uninstalled if they are currently being used. Features turned on by permanent licenses must first be disabled, before that license is uninstalled.



Tip

If you are using an evaluation license and would like to install a new permanent license, you can do so without service disruption and before the evaluation license expires. Removing an evaluation license immediately triggers a grace period without service disruption.



Caution

Uninstalling a license requires the related features to first be disabled.

Send documentation comments to mdsfeedback-doc@cisco.com.

To uninstall a license, follow these steps:

-
- Step 1** Save your running configuration to a remote server using the **copy** command (see the “[Initial Configuration](#)” section on page 4-1).
- Step 2** Issue the **show license brief** command in EXEC mode to view a list of all installed license key files and identify the file to be uninstalled. In this example, the file to be uninstalled is the Ficon.lic file.
- ```
switch# show license brief
Enterprise.lic
Ficon.lic
```
- Step 3** Disable the features provided by the license to be uninstalled. Issue the **show license usage package\_name** command to view the enabled features for a specified package.
- ```
switch# show license usage ENTERPRISE_PKG
Application
-----
ivr
qos_manager
-----
```
- Step 4** Uninstall the Ficon.lic file using the **clear license filename** command, where *filename* is the name of the installed license key file.
- ```
switch# clear license Enterprise.lic
Clearing license Enterprise.lic:
SERVER this_host ANY
VENDOR cisco
```
- Step 5** Enter **yes** (yes is the default) to continue with the license update.
- ```
Do you want to continue? (y/n) y
Clearing license ..done
```
- The Ficon.lic license key file is now uninstalled.
-

Updating Licenses

If your license is time bound, you must obtain and install an updated license. Contact technical support to request an updated license.



Note

If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco Systems, contact Cisco Technical Support at this URL: <http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

To update a license, follow these steps:

-
- Step 1** Obtain the updated license file using the procedure described in the “[Obtaining the License Key File](#)” section on page 3-7.
- Step 2** Save your running configuration to a remote server using the **copy** command (see the “[Copying Files](#)” section on page 4-28).
- Step 3** Issue the **show license brief** command to verify the name of the file to be updated.

Send documentation comments to mdsfeedback-doc@cisco.com.

```
switch# show license brief
sanextn1.lic:
```

- Step 4** Update the license file using the **update license url** command, where *url* specifies the bootflash:, slot0:, or volatile: directory location of the updated license file.

```
switch# update license bootflash:sanextn2.lic sanextn1.lic
Updating sanextn1.lic:
SERVER this_host ANY
VENDOR cisco
# An example fcports license
INCREMENT SAN_EXTN_OVER_IP cisco 1.000 permanent 1 HOSTID=VDH=ABCD \
    NOTICE=<LicFileID>san_extn1.lic</LicFileID><LicLineID>0</LicLineID> \
    SIGN=33088E76F668

with bootflash:/sanextn2.lic:
SERVER this_host ANY
VENDOR cisco
# An example fcports license
INCREMENT SAN_EXTN_OVER_IP cisco 1.000 permanent 1 HOSTID=VDH=ABCD \
    NOTICE=<LicFileID>san_extn2.lic</LicFileID><LicLineID>1</LicLineID> \
    SIGN=67CB2A8CCAC2
```

- Step 5** Enter **yes** (yes is the default), to continue with the license update.

```
Do you want to continue? (y/n) y
Updating license ..done
switch#
The sanextn1.lic license key file is now updated.
```

License Expiry Alerts

The Cisco SAN-OS license counter keeps track of all licenses on a switch. Once an expiry period has started, you will receive console messages, SNMP traps, system messages, and Call Home messages on a daily basis.

Beyond that, the frequency of these messages become hourly during the last seven days of the expiry time span. The following example uses the FICON license feature.

Your FICON license feature is scheduled to expire in 60 days. If today is December 1st, the license expires on January 30th. In this case, you will receive:

- Daily alerts from December 1st to January 23rd.
- Hourly alerts from January 24th to January 29th.
- From January 30th, the FICON feature runs without a license for a grace period of 60 days (for earlier versions of SAN-OS 1.3(3c) or earlier) or 120 days (for SAN-OS 1.3(4) or later).
- From January 30th to May 21st, you receive daily alerts about the grace period usage.
- From May 22nd to May 30th, you receive hourly alerts about the grace period ending.
- On May 31st, the FICON feature is automatically turned off.



Note

License expiry alerts cannot be configured.

Send documentation comments to mdsfeedback-doc@cisco.com.



Caution

After the final seven days of the grace period, the feature is turned off and your network traffic may be disrupted. The grace period also applies to licensed features in Cisco MDS SAN-OS Release 1.2. While Cisco MDS SAN-OS Release 1.2 did not enforce licenses, any future upgrade will enforce license requirements and the 60-day grace period (for earlier versions of SAN-OS 1.3(3c) or earlier) or 120-day grace period (for SAN-OS 1.3(4) or later).

Grace Period Countdown

The grace period is set to 60 days (for earlier versions of SAN-OS 1.3(3c) or earlier) or 120 days (for SAN-OS 1.3(4) or later), and the countdown starts or continues when either of the following two situations occur:

- You are evaluating a feature for which you have not purchased a license.
- You purchased a license that has reached its expiry date and you are still using the feature requiring that license.

License packages can contain several features. If you disable a feature during the grace period and there are other features in that license package which are still enabled, the clock does not stop for that license package. To suspend the grace period countdown for a licensed feature, you must disable every feature in that license package.

License Transfers Between Switches

A license is specific to the switch for which it is issued and is not valid on any other switch. If you need to transfer a license from one switch to another, contact your customer service representative.



Note

If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco Systems, contact Cisco Technical Support at this URL:
<http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

Displaying License Information

Use the **show license** commands to display all license information configured on this switch (see Examples 3-1 to 3-6).

Example 3-1 Displays Information About Current License Usage

```
switch# show license usage
```

| Feature | Installed | License Count | Status | ExpiryDate | Comments |
|-----------------------|-----------|---------------|--------|------------|--------------------------|
| FM_SERVER_PKG | Yes | - | Unused | never | license missing |
| MAINFRAME_PKG | No | - | Unused | | Grace Period 57days15hrs |
| ENTERPRISE_PKG | Yes | - | InUse | never | - |
| SAN_EXTN_OVER_IP | No | 0 | Unused | | - |
| SAN_EXTN_OVER_IP_IPS4 | No | 0 | Unused | | - |

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 3-2 Displays the List of Features in a Specified Package

```
switch# show license usage ENTERPRISE_PKG
Application
-----
ivr
qos_manager
-----
```

Example 3-3 Displays the Host ID for the License

```
switch# show license host-id
License hostid: VDH=FOX0646S017
```



Note

Use the entire ID that appears after the colon (:) sign. The VHD is the Vendor Host ID.

Example 3-4 Displays All Installed License Key Files and their Contents

```
switch# show license
Permanent.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT MAINFRAME_PKG cisco 1.0 permanent uncounted \
    HOSTID=VDH=FOX0646S017 \
    NOTICE="<LicFileID></LicFileID><LicLineID>0</LicLineID> \
    <PAK>dummyPak</PAK>" SIGN=EE9F91EA4B64
Evaluation.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT MAINFRAME_PKG cisco 1.0 30-Dec-2003 uncounted \
    HOSTID=VDH=FOX0646S017 \
    NOTICE="<LicFileID></LicFileID><LicLineID>0</LicLineID> \
    <PAK>dummyPak</PAK>" SIGN=EE9F91EA4B64
```

Example 3-5 Displays a List of Installed License Key Files

```
switch# show license brief
Enterprise.lic
Ficon.lic
FCIP.lic
```

Example 3-6 Displays the Contents of a Specified License Key File

```
switch# show license file Permanent.lic
Permanent.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT MAINFRAME_PKG cisco 1.0 permanent uncounted \
    HOSTID=VDH=FOX0646S017 \
    NOTICE="<LicFileID></LicFileID><LicLineID>0</LicLineID> \
    <PAK>dummyPak</PAK>" SIGN=EE9F91EA4B64
```

Send documentation comments to mdsfeedback-doc@cisco.com.



Initial Configuration

This chapter describes how to initially configure switches so they can be accessed by other devices. This chapter includes the following sections:


- [Starting a Switch in the Cisco MDS 9000 Family, page 4-2](#)
- [Initial Setup Routine, page 4-2](#)
- [Accessing the Switch, page 4-14](#)
- [Assigning a Switch Name, page 4-14](#)
- [Where Do You Go Next?, page 4-15](#)
- [Verifying the Module Status, page 4-15](#)
- [Configuring Date and Time, page 4-16](#)
- [Management Interface Configuration, page 4-21](#)
- [Telnet Server Connection, page 4-24](#)
- [Working with Configuration Files, page 4-24](#)
- [Downgrading from a Higher Release, page 4-31](#)
- [Accessing Remote File Systems, page 4-33](#)
- [Deleting Files, page 4-33](#)
- [Configuring Console Port Settings, page 4-34](#)
- [Configuring COM1 Port Settings, page 4-35](#)
- [Configuring Modem Connections, page 4-36](#)
- [Configuring CDP, page 4-40](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

Starting a Switch in the Cisco MDS 9000 Family

The following procedure is a review of the tasks you should have completed during hardware installation, including starting up the switch. These tasks must be completed before you can configure the switch.

Before you can configure a switch, follow these steps:

-
- Step 1** Verify the following physical connections for the new Cisco MDS 9000 Family switch:
- The console port is physically connected to a computer terminal (or terminal server).
 - The management 10/100 Ethernet port (mgmt0) is connected to an external hub, switch, or router.
- Refer to the *Cisco MDS 9000 Family Hardware Installation Guide* (for the required product) for more information.
-  **Tip** Save the host ID information for future use (for example, to enable licensed features). The host ID information is provided in the Proof of Purchase document that accompanies the switch.
-
- Step 2** Verify that the default console port parameters are identical to those of the computer terminal (or terminal server) attached to the switch console port:
- 9600 baud
 - 8 data bits
 - 1 stop bit
 - No parity
- Step 3** Power on the switch. The switch boots automatically and the `switch#` prompt appears in your terminal window.
-

Initial Setup Routine

The first time that you access a switch in the Cisco MDS 9000 Family, it runs a setup program that prompts you for the IP address and other configuration information necessary for the switch to communicate over the supervisor module Ethernet interface. This information is required to configure and manage the switch.



Note

The IP address can only be configured from the CLI. When you power up the switch for the first time assign the IP address. After you perform this step, the Cisco MDS 9000 Family Fabric Manager can reach the switch through the console port.

Send documentation comments to mdsfeedback-doc@cisco.com.

Preparing to Configure the Switch

Before you configure a switch in the Cisco MDS 9000 Family for the first time, you need the following information:

- Administrator password, including:
 - Creating a password for the administrator (required).
 - Creating an additional login account and password (optional).
- IP address for the switch management interface—The management interface can be an out-of-band Ethernet interface or an in-band Fibre Channel interface (recommended).
- Subnet mask for the switch's management interface (optional).
- IP addresses, including:
 - Destination prefix, destination prefix subnet mask, and next hop IP address, if you want to enable IP routing. Also, provide the IP address of the default network (optional).
 - Otherwise, provide an IP address of the default gateway (optional).
- SSH service on the switch—To enable this optional service, select the type of SSH key (dsa/rsa/rsa1) and number of key bits (768 to 2048).
- DNS IP address (optional).
- Default domain name (optional).
- NTP server IP address (optional).
- SNMP community string (optional).
- Switch name—This is your switch prompt (optional).



Note

Be sure to configure the IP route, the IP default network address, and the IP default gateway address to enable SNMP access. If IP routing is enabled, the switch uses the IP route and the default network IP address. If IP routing is disabled, the switch uses the default gateway IP address.

Default Login

All Cisco MDS 9000 Family switches have the network administrator as a default user (admin). You cannot change the default user at any time (see the [“Role-Based Authorization” section on page 19-21](#)).

As of Release 2.0(1b), you must explicitly configure a strong password for any switch in the Cisco MDS 9000 Family. If a password is trivial (short, easy-to-decipher), your password configuration is rejected. Be sure to configure a strong password (see [“Configuring User Accounts” section on page 19-29](#)). If you configure and subsequently forget this new password, you have the option to recover this password (see the [“Recovering Administrator Password” section on page 19-37](#)).

Send documentation comments to mdsfeedback-doc@cisco.com.

Setup Options

The setup scenario differs based on the subnet to which you are adding the new switch. You must configure a Cisco MDS 9000 Family switch with an IP address to enable management connections from outside of the switch.

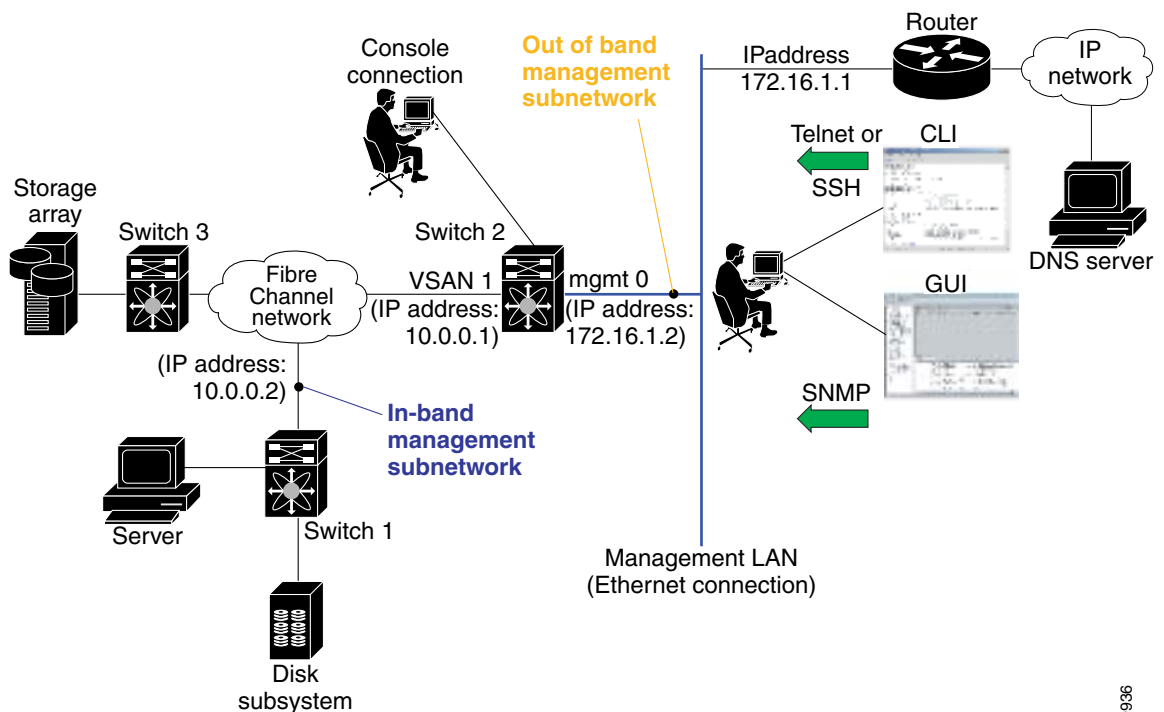


Note

Some concepts such as out-of-band management and in-band management are briefly explained here. These concepts are explained in more detail in subsequent chapters.

- Out-of-band management—This feature provides a connection to the network through a supervisor module front panel Ethernet port (see [Figure 4-1](#)).
- In-band management—This feature provides IP over Fibre Channel (IPFC) to manage the switches. The in-band management feature is transparent to the network management system (NMS). Instead of conventional Ethernet physical media, switches in the Cisco MDS 9000 Family use IPFC as the transport mechanism (see [Figure 4-1](#) and [Chapter 26, “Configuring IP Services”](#)).

Figure 4-1 Management Access to Switches



79936

Send documentation comments to mdsfeedback-doc@cisco.com.

Assigning Setup Information

This section describes how to initially configure the switch for both out-of-band and in-band management.



Note

Press **Ctrl-C** at any prompt to skip the remaining configuration options and proceed with what is configured until that point. Entering the new password for the administrator is a requirement and cannot be skipped.



Tip

If you do not wish to answer a previously configured question, or if you wish to skip answers to any questions, press **Enter**. If a default answer is not available (for example, switch name), the switch uses what was previously configured and skips to the next question.

Configuring Out-of-Band Management



Note

You can configure both in-band and out-of-band configuration together by entering **Yes** in both [Step 11c.](#) and [Step 11d.](#) in the following procedure.

To configure the switch for first time out-of-band access, follow these steps:

Step 1 Power on the switch. Switches in the Cisco MDS 9000 Family boot automatically.

Step 2 Enter the new password for the administrator.

Enter the password for admin: **2004asdf*1kjh18**



Tip

If a password is trivial (short, easy-to-decipher), your password configuration is rejected. Be sure to configure a strong password as shown in the sample configuration. Passwords are case-sensitive. As of Cisco SAN-OS Release 2.0(1b), **admin** is not the default password for any switch in the Cisco MDS 9000 Family. You must explicitly configure a password that meets the requirements listed in this tip (see [“Configuring User Accounts”](#) section on page 19-29).

Step 3 Enter **yes** to enter the setup mode.

This setup utility will guide you through the basic configuration of the system. Setup configures only enough connectivity for management of the system.

Please register Cisco MDS 9000 Family devices promptly with your supplier. Failure to register may affect response times for initial service calls. MDS devices must be registered to receive entitled support services.

Press Enter incase you want to skip any dialog. Use ctrl-c at anytime to skip away remaining dialogs.

Would you like to enter the basic configuration dialog (yes/no): **yes**

The setup utility guides you through the basic configuration process. Press **Ctrl-C** at any prompt to end the configuration process.

Send documentation comments to mdsfeedback-doc@cisco.com.

- Step 4** Enter the new password for the administrator (admin is the default).

Enter the password for admin: **admin**

- Step 5** Enter **yes** (no is the default) to create additional accounts.

Create another login account (yes/no) [n]: **yes**

While configuring your initial setup, you can create an additional user account (in the network-admin role) besides the administrator's account. See the [“Role-Based Authorization”](#) section on page 19-21 for information on default roles and permissions.



Note User login IDs must contain non-numeric characters.

- a. Enter the user login ID.

Enter the user login ID: *user_name*

- b. Enter the user password.

Enter the password for user_name: *user-password*

- Step 6** Enter **yes** (yes is the default) to create an SNMPv3 account.

Configure SNMPv3 Management parameters (yes/no) [y]: **yes**

- a. Enter the user name (admin is the default).

SNMPv3 user name [admin]: **admin**

- b. Enter the SNMPv3 password (minimum of eight characters). The default is **admin123**.

SNMPv3 user authentication password: *admin_pass*



Note By default, if the admin password is at least eight characters, then the SNMP authentication password is the same as the admin password (at least eight characters). If the admin password is less than eight characters, then you need to provide a new password for SNMP. The admin password can have a minimum of one character, but the SNMP authentication password must have a minimum of eight characters.

- Step 7** Enter **yes** (no is the default) to configure the read-only or read-write SNMP community string.

Configure read-only SNMP community string (yes/no) [n]: **yes**

- a. Enter the SNMP community string.

SNMP community string: *snmp_community*

- Step 8** Enter a name for the switch.



Note The switch name is limited to 32 alphanumeric characters. The default is **switch**.

Enter the switch name: *switch_name*

Send documentation comments to mdsfeedback-doc@cisco.com.

Step 9 Enter **yes** (yes is the default) to configure out-of-band management.

Continue with Out-of-band (mgmt0) management configuration? [yes/no]: **yes**

a. Enter the mgmt0 IP address.

Mgmt0 IP address: *ip_address*

b. Enter the mgmt0 subnet mask.

Mgmt0 IP netmask: *subnet_mask*

Step 10 Enter **yes** (yes is the default) to configure the default gateway (recommended).

Configure the default-gateway: (yes/no) [y]: **yes**

a. Enter the default gateway IP address.

IP address of the default-gateway: *default_gateway*

Step 11 Enter **yes** (**no** is the default) to configure advanced IP options such as inband management, static routes, default network, DNS and domain name.

Configure Advanced IP options (yes/no)? [n]: **yes**

a. Enter **no** (no is the default) at the in-band management configuration prompt.

Continue with in-band (VSAN1) management configuration? (yes/no) [no]: **no**

b. Enter **yes** (yes is the default) to enable IP routing capabilities.

Enable the ip routing? (yes/no) [y]: **yes**

c. Enter **yes** (yes is the default) to configure a static route (recommended).

Configure static route: (yes/no) [y]: **yes**

Enter the destination prefix.

Destination prefix: *dest_prefix*

Type the destination prefix mask.

Destination prefix mask: *dest_mask*

Type the next hop IP address.

Next hop ip address: *next_hop_address*



Note

Be sure to configure the IP route, the default network IP address, and the default gateway IP address to enable SNMP access. If IP routing is enabled, the switch uses the IP route and the default network IP address. If IP routing is disabled, the switch uses the default gateway IP address.

d. Enter **yes** (yes is the default) to configure the default network (recommended).

Configure the default network: (yes/no) [y]: **yes**

Enter the default network IP address.

Send documentation comments to mdsfeedback-doc@cisco.com.



Note The default network IP address is the destination prefix provided in [Step 11c](#).

Default network IP address [dest_prefix]: *dest_prefix*

- e. Enter **yes** (yes is the default) to configure the DNS IP address.

Configure the DNS IP address? (yes/no) [y]: **yes**

Enter the DNS IP address.

DNS IP address: *name_server*

- f. Enter **yes** (default is no) to configure the default domain name.

Configure the default domain name? (yes/no) [n]: **yes**

Enter the default domain name.

Default domain name: *domain_name*

- Step 12** Enter **yes** (yes is the default) to enable Telnet service.

Enable the telnet service? (yes/no) [y]: **yes**

- Step 13** Enter **yes** (no is the default) to enable the SSH service.

Enabled SSH service? (yes/no) [n]: **yes**

- Step 14** Enter the SSH key type (see the “[Generating the SSH Server Key Pair](#)” section on page 19-35) that you would like to generate.

Type the SSH key you would like to generate (dsa/rsa/rsa1)? **dsa**

- Step 15** Enter the number of key bits within the specified range.

Enter the number of key bits? (768 to 2048): **768**

- Step 16** Enter **yes** (no is the default) to configure the NTP server.

Configure NTP server? (yes/no) [n]: **yes**

- a. Enter the NTP server IP address.

NTP server IP address: *ntp_server_IP_address*

- Step 17** Enter **shut** (shut is the default) to configure the default switchport interface to the shut state.

Configure default switchport interface state (shut/noshut) [shut]: **shut**



Note The management ethernet interface is not shut down at this point—only the Fibre Channel, iSCSI, FCIP, and Gigabit Ethernet interfaces are shut down.

- Step 18** Enter **on** (on is the default) to configure the switchport trunk mode.

Configure default switchport trunk mode (on/off/auto) [on]: **on**

- Step 19** Enter **on** (off is the default) to configure the PortChannel auto-create state.

Configure default port-channel auto-create state (on/off) [off]: **on**

Send documentation comments to mdsfeedback-doc@cisco.com.

Step 20 Enter **permit** (deny is the default) to deny a default zone policy configuration.

Configure default zone policy (permit/deny) [deny]: **permit**

Permits traffic flow to all members of the default zone.

Step 21 Enter **yes** (no is the default) to disable a full zone set distribution (see the “Zone Set Distribution” section on page 15-11).

Enable full zoneset distribution (yes/no) [n]: **yes**

Disables the switch-wide default for the full zone set distribution feature.

You see the new configuration. Review and edit the configuration that you have just entered.

Step 22 Enter **no** (no is the default) if you are satisfied with the configuration.

The following configuration will be applied:

```
username admin password admin_pass role network-admin
username user_name password user_pass role network-admin
snmp-server community snmp_community ro
switchname switch
interface mgmt0
    ip address ip_address subnet_mask
    no shutdown
ip routing
ip route dest_prefix dest_mask dest_address
ip default-network dest_prefix
ip default-gateway default_gateway
ip name-server name_server
ip domain-name domain_name
telnet server enable
ssh key dsa 768 force
ssh server enable
ntp server ipaddr ntp_server
system default switchport shutdown
system default switchport trunk mode on
system default port-channel auto-create
zone default-zone permit vsan 1-4093
zoneset distribute full vsan 1-4093
```

Would you like to edit the configuration? (yes/no) [n]: **no**

Step 23 Enter **yes** (yes is default) to use and save this configuration:

Use this configuration and save it? (yes/no) [y]: **yes**



Caution

If you do not save the configuration at this point, none of your changes are updated the next time the switch is rebooted. Type **yes** in order to save the new configuration. This ensures that the kickstart and system images are also automatically configured (see [Chapter 6, “Software Images”](#)).

Configuring In-Band Management

The in-band management logical interface is VSAN 1. This management interface uses the Fibre Channel infrastructure to transport IP traffic. An interface for VSAN 1 is created on every switch in the fabric. Each switch should have its VSAN 1 interface configured with an IP address in the same

Send documentation comments to mdsfeedback-doc@cisco.com.

subnetwork. A default route that points to the switch providing access to the IP network should be configured on every switch in the Fibre Channel fabric (see [Chapter 10, “Configuring and Managing VSANs”](#)).



Note

You can configure both in-band and out-of-band configuration together by entering **Yes** in both [Step 9c](#) and [Step 9d](#) in the following procedure.

To configure a switch for first time in-band access, follow these steps:

Step 1 Power on the switch. Switches in the Cisco MDS 9000 Family boot automatically.

Step 2 Enter the new password for the administrator.

Enter the password for admin: **2004asdf*1kj18**



Tip

If a password is trivial (short, easy-to-decipher), your password configuration is rejected. Be sure to configure a strong password as shown in the sample configuration. Passwords are case-sensitive. As of Cisco SAN-OS Release 2.0(1b), **admin** is not the default password for any switch in the Cisco MDS 9000 Family. You must explicitly configure a password that meets the requirements listed in this tip (see [“Configuring User Accounts” section on page 19-29](#)).

Step 3 Enter **yes** to enter the setup mode.

This setup utility will guide you through the basic configuration of the system. Setup configures only enough connectivity for management of the system.

Please register Cisco MDS 9000 Family devices promptly with your supplier. Failure to register may affect response times for initial service calls. MDS devices must be registered to receive entitled support services.

Press Enter incase you want to skip any dialog. Use ctrl-c at anytime to skip away remaining dialogs.

Would you like to enter the basic configuration dialog (yes/no): **yes**

The setup utility guides you through the basic configuration process. Press **Ctrl-C** at any prompt to end the configuration process.

Step 4 Enter **no** (no is the default) if you do not wish to create additional accounts.

Create another login account (yes/no) [no]: **no**

Step 5 Configure the read-only or read-write SNMP community string.

a. Enter **no** (no is the default) to avoid configuring the read-only SNMP community string.

Configure read-only SNMP community string (yes/no) [n]: **no**

b. Enter **no** (no is the default) to configure the read-only SNMP community string.

Configure read-only SNMP community string (yes/no) [n]: **yes**

c. Enter the SNMP community string.

SNMP community string: *snmp_community*

Step 6 Enter a name for the switch.

Send documentation comments to mdsfeedback-doc@cisco.com.



Note The switch name is limited to 32 alphanumeric characters. The default is **switch**.

Enter the switch name: *switch_name*

Step 7 Enter **no** (yes is the default) at the configuration prompt to configure out-of-band management.

Continue with Out-of-band (mgmt0) management configuration? [yes/no]: **no**

Step 8 Enter **yes** (yes is the default) to configure the default gateway.

Configure the default-gateway: (yes/no) [y]: **yes**

a. Enter the default gateway IP address.

IP address of the default gateway: *default_gateway*

Step 9 Enter **yes** (**no** is the default) to configure advanced IP options such as Inband management, static routes, default network, dns and domain name.

Configure Advanced IP options (yes/no)? [n]: **yes**

a. Enter **yes** (no is the default) at the in-band management configuration prompt.

Continue with in-band (VSAN1) management configuration? (yes/no) [no]: **yes**

Enter the VSAN 1 IP address.

VSAN1 IP address: *ip_address*

Enter the subnet mask.

VSAN1 IP net mask: *subnet_mask*

b. Enter **no** (yes is the default) to enable IP routing capabilities.

Enable ip routing capabilities? (yes/no) [y]: **no**

c. Enter **no** (yes is the default) to configure a static route.

Configure static route: (yes/no) [y]: **no**

d. Enter **no** (yes is the default) to configure the default network.

Configure the default-network: (yes/no) [y]: **no**

e. Enter **no** (yes is the default) to configure the DNS IP address.

Configure the DNS IP address? (yes/no) [y]: **no**

f. Enter **no** (no is the default) to skip the default domain name configuration.

Configure the default domain name? (yes/no) [n]: **no**

Step 10 Enter **no** (yes is the default) to disable Telnet service.

Enable the telnet service? (yes/no) [y]: **no**

Step 11 Enter **yes** (no is the default) to enable the SSH service.

Enabled SSH service? (yes/no) [n]: **yes**

Send documentation comments to mdsfeedback-doc@cisco.com.

- Step 12** Enter the SSH key type (see the “[Generating the SSH Server Key Pair](#)” section on page 19-35) that you would like to generate.

Type the SSH key you would like to generate (dsa/rsa/rsal)? **rsa**

- Step 13** Enter the number of key bits within the specified range.

Enter the number of key bits? (768 to 1024): **1024**

- Step 14** Enter **no** (no is the default) to configure the NTP server.

Configure NTP server? (yes/no) [n]: **no**

- Step 15** Enter **shut** (shut is the default) to configure the default switchport interface to the shut state.

Configure default switchport interface state (shut/noshut) [shut]: **shut**



Note The management Ethernet interface is not shut down at this point—only the Fibre Channel, iSCSI, FCIP, and Gigabit Ethernet interfaces are shut down.

- Step 16** Enter **auto** (off is the default) to configure the switchport trunk mode.

Configure default switchport trunk mode (on/off/auto) [off]: **auto**

- Step 17** Enter **off** (off is the default) to configure the PortChannel auto-create state.

Configure default port-channel auto-create state (on/off) [off]: **off**

- Step 18** Enter **deny** (deny is the default) to deny a default zone policy configuration.

Configure default zone policy (permit/deny) [deny]: **deny**

Denies traffic flow to all members of the default zone.

- Step 19** Enter **no** (no is the default) to disable a full zone set distribution (see the “[Zone Set Distribution](#)” section on page 15-11).

Enable full zoneset distribution (yes/no) [n]: **no**

Disables the switch-wide default for the full zone set distribution feature.

You see the new configuration. Review and edit the configuration that you have just entered.

Send documentation comments to mdsfeedback-doc@cisco.com.

Step 20 Enter **no** (no is the default) if you are satisfied with the configuration.

The following configuration will be applied:

```
username admin password admin_pass role network-admin
snmp-server community snmp_community rw
switchname switch
interface vsan1
    ip address ip_address subnet_mask
    no shutdown
ip default-gateway default_gateway
no telnet server enable
ssh key rsa 1024 force
ssh server enable
no system default switchport shutdown
system default switchport trunk mode auto
no zone default-zone permit vsan 1-4093
no zoneset distribute full vsan 1-4093
```

Would you like to edit the configuration? (yes/no) [n]: **no**

Step 21 Enter **yes** (yes is default) to use and save this configuration.

Use this configuration and save it? (yes/no) [y]: **yes**



Caution

If you do not save the configuration at this point, none of your changes are updated the next time the switch is rebooted. Type **yes** in order to save the new configuration. This ensures that the kickstart and system images are also automatically configured (see [Chapter 6, “Software Images”](#)).

Using the setup Command

To make changes to the initial configuration at a later time, you can issue the **setup** command in EXEC mode.

```
switch# setup
---- Basic System Configuration Dialog ----
This setup utility will guide you through the basic configuration of
the system. Setup configures only enough connectivity for management
of the system.
*Note: setup always assumes a predefined defaults irrespective
of the current system configuration when invoked from CLI.
```

Press Enter incase you want to skip any dialog. Use ctrl-c at anytime to skip away remaining dialogs.

Would you like to enter the basic configuration dialog (yes/no): **yes**

The setup utility guides you through the basic configuration process.

Send documentation comments to mdsfeedback-doc@cisco.com.

Accessing the Switch

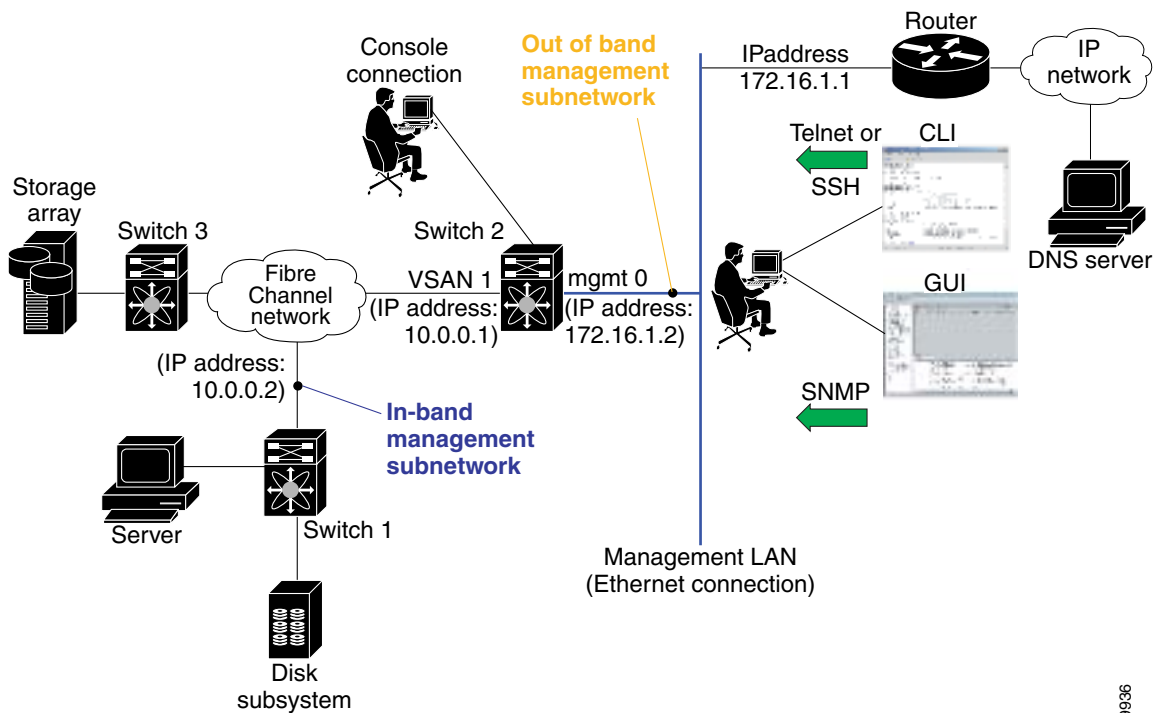
After initial configuration, you can access the switch in one of three ways (see [Figure 4-2](#)):

- Serial console access—You can use a serial port connection to access the CLI.
- In-band IP (IPFC) access—You can use Telnet or SSH to access a switch in the Cisco MDS 9000 Family or use SNMP to connect to a Cisco MDS 9000 Fabric Manager application.
- Out-of-band (10/100BASE-T Ethernet) access—You can use Telnet or SSH to access a switch in the Cisco MDS 9000 Family or use SNMP to connect to a Cisco MDS 9000 Fabric Manager application.



Note To use the Cisco Fabric Manager, refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*.

Figure 4-2 Switch Access Options



79936

Assigning a Switch Name

Each switch in the fabric requires a unique name. You can assign names to easily identify the switch by its physical location, its SAN association, or the organization to which it is deployed. The assigned name is displayed in the command-line prompt. The switch name is limited to 20 alphanumeric characters.



Note This guide refers to a switch in the Cisco MDS 9000 Family as *switch*, and it uses the `switch#` prompt.

Send documentation comments to mdsfeedback-doc@cisco.com.

To change the name of the switch, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# switchname myswitch1 myswitch1(config)# | Changes the switch name prompt as specified. |
| Step 3 | myswitch1(config)# no switchname switch(config)# | Reverts the switch name prompt to its default (switch#). |

Where Do You Go Next?

After reviewing the default configuration, you can change it or perform other configuration or management tasks. The initial setup can only be performed at the CLI. However, you can continue to configure other software features, or access the switch after initial configuration by using either the CLI or the Device Manager and Fabric Manager applications.

To use the Cisco MDS 9000 Fabric Manager, refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*.

Verifying the Module Status

Before you begin configuring the switch, you need to ensure that the modules in the chassis are functioning as designed. To verify the status of a module at any time, issue the **show module** command in EXEC mode. A sample output of the **show module** command follows:

```
switch# show module
Mod  Ports  Module-Type                Model                Status
---  ---
2    8      IP Storage Services Module DS-X9308-SMIP        ok
5    0      Supervisor/Fabric-1       DS-X9530-SF1-K9     active *
6    0      Supervisor/Fabric-1       DS-X9530-SF1-K9     ha-standby
8    0      Caching Services Module   DS-X9560-SMAP        ok
9    32     1/2 Gbps FC Module        DS-X9032              ok

Mod  Sw          Hw          World-Wide-Name(s) (WWN)
---  ---
2    1.3(0.106a) 0.206       20:41:00:05:30:00:00:00 to 20:48:00:05:30:00:00:00
5    1.3(0.106a) 0.602       --
6    1.3(0.106a) 0.602       --
8    1.3(0.106a) 0.702       --
9    1.3(0.106a) 0.3         22:01:00:05:30:00:00:00 to 22:20:00:05:30:00:00:00

Mod  MAC-Address(es)                Serial-Num
---  ---
2    00-05-30-00-9d-d2 to 00-05-30-00-9d-de JAB064605a2
5    00-05-30-00-64-be to 00-05-30-00-64-c2 JAB06350B1R
6    00-d0-97-38-b3-f9 to 00-d0-97-38-b3-fd JAB06350B1R
8    00-05-30-01-37-7a to 00-05-30-01-37-fe JAB072705ja
9    00-05-30-00-2d-e2 to 00-05-30-00-2d-e6 JAB06280ae9
```

* this terminal session

If the status is OK or active, you can continue with your configuration (see [Chapter 7, “Managing Modules”](#)).

Send documentation comments to mdsfeedback-doc@cisco.com.

Configuring Date and Time

Switches in the Cisco MDS 9000 Family use Universal Coordinated Time (UTC), which is the same as Greenwich Mean Time (GMT). To change the default time on the switch, issue the **clock** command from EXEC mode.

```
switch# clock set <HH:MM:SS> <DD> <Month in words> <YYYY>
```

For example:

```
switch# clock set 15:58:09 23 September 2002
Mon Sep 23 15:58:09 UTC 2002
```

Where *HH* represents hours in military format (15 for 3 p.m.), *MM* is minutes (58), *SS* is seconds (09), *DD* is the date (23), *Month* is the month in words (September), and *YYYY* is the year (2002).



Note

The **clock** command changes are saved across system resets.

Configuring the Time Zone

You can specify a time zone for the switch.

To specify the local time without the daylight savings feature, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# confi g t | Enters configuration mode. |
| Step 2 | switch(config)# clock timezone <timezone name> <-23 to 23 hours offset from UTC time> <0 to 50 minutes offset from UTC> Example: switch(config)# clock timezone PST -8 0 | Sets the time zone with a specified name, specified hours, and specified minutes. This example sets the time zone to Pacific Standard Time (PST) and offsets the UTC time by negative eight hours and 0 minutes. |
| Step 3 | switch(config)# exit switch# | Returns to EXEC mode. |
| Step 4 | switch# show clock | Verifies the time zone configuration. |
| Step 5 | switch# show run | Displays changes made to the time zone configuration along with other configuration information. |

Adjusting for Daylight Saving Time

Following U.S. standards, you can have the switch advance the clock one hour at 2:00 a.m. on the first Sunday in April and move back the clock one hour at 2:00 a.m. on the last Sunday in October. You can also explicitly specify the start and end dates and times and whether or not the time adjustment recurs every year.

Send documentation comments to mdsfeedback-doc@cisco.com.

To enable the daylight saving time clock adjustment according to the U.S. rules, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# clock timezone <i>timezone_name hour_offset_from_UTC</i> <i>minute_offset_from_UTC</i> Example: switch(config)# clock timezone PST -8 0 | Offsets the time zone as specified. This example set the Pacific standard offset time as negative 8 hours and 0 minutes. |
| Step 3 | switch(config)# no clock timezone switch(config)# clock summer-time <i>daylight_timezone_name start_week</i> <i>start_day start_month start_time end_week</i> <i>end_day end_month end_time</i> <i>daylight_offset_inminutes</i> Example: switch(config)# clock summer-time PDT 1 Sun Apr 02:00 5 Sun Oct 02:00 60 switch(config)# | Disables the time zone adjustment feature. Sets the daylight savings time for a specified time zone. The start and end values are as follows: <ul style="list-style-type: none"> • Week ranging from 1 through 5 • Day ranging from Sunday through Saturday • Month ranging from January through December The daylight offset ranges from 1 through 1440 minutes which are added to the start time and deleted time from the end time. This example adjusts the daylight savings time for the Pacific daylight time by 60 minutes starting the first Sunday in April at 2 a.m. and ending the last Sunday in October at 2 a.m. |
| | switch(config)# no clock summer-time | Disables the daylight saving time adjustment feature. |
| Step 4 | switch(config)# exit switch# | Returns to EXEC mode. |
| Step 5 | switch# show clock | Verifies the time zone configuration. |

NTP Configuration

A Network Time Protocol (NTP) server provides a precise time source (radio clock or atomic clock) to synchronize the system clocks of network devices. NTP is transported over User Datagram Protocol UDP/IP. All NTP communications use UTC. An NTP server receives its time from a reference time source, such as a radio clock or atomic clock, attached to the time. NTP distributes this time across the network.

In a large enterprise network, having one time standard for all network devices is critical for management reporting and event logging functions when trying to correlate interacting events logged across multiple devices. Many enterprise customers with extremely mission-critical networks maintain their own stratum-1 NTP source.

Time synchronization happens when several frames are exchanged between clients and servers. The switches in client mode know the address of one or more NTP servers. The servers act as the time source and receive client synchronization requests.

Send documentation comments to mdsfeedback-doc@cisco.com.

By configuring an IP address as a peer, the switch will obtain and provide time as required. The peer is capable of providing time on its own and is capable of having a server configured. If both these instances point to different time servers, your NTP service is more reliable. Thus, even if the active server link is lost, you can still maintain the right time due to the presence of the peer.


Tip

If an active server fails, a configured peer helps in providing the NTP time. Provide a direct NTP server association and configure a peer to ensure backup support if the active server fails.

If you only configure a peer, the most accurate peer takes on the role of the NTP server and the other peer(s) act as a peer(s). Both machines end at the right time if they have the right time source or if they point to the right NTP source.

To configure NTP in a server association, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# conf t | Enters configuration mode. |
| Step 2 | switch(config)# ntp server 10.10.10.10 switch(config)# | Forms a server association with a server. |
| Step 3 | switch(config)# ntp peer 10.20.10.0 switch(config)# | Forms a peer association with a peer. You can specify multiple associations. |
| Step 4 | switch(config)# exit switch# | Returns to EXEC mode. |
| Step 5 | switch# copy running-config startup-config | Saves your configuration changes to NVRAM. |
| Step 6 | switch# show ntp peers ----- Peer IP Address Serv/Peer ----- 10.20.10.2 Server 10.20.10.0 Peer | Tip This is one instance where you can save the configuration as a result of an NTP configuration change. You can issue this command at any time. |
| | | Note A domain name is resolved only when you have a DNS server configured. |

NTP Configuration Guidelines

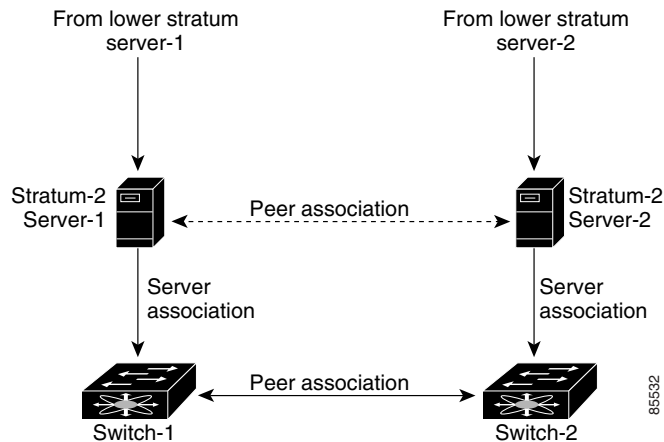
The following guidelines apply to all NTP configurations:

- You should have a peer association with another switch only when you are sure that your clock is reliable (which means that you are a client of a reliable NTP server).
- A peer configured alone takes on the role of a server and should be used as backup. If you have two servers, then you can have several switches point to one server, and the remaining switches to the other server. Then you would configure peer association between these two sets. This forces the clock to be more reliable.
- If you only have one server, it's better for all the switches to have a client association with that server.

Not even a server down time will affect well-configured switches in the network. [Figure 4-3](#) displays a network with two NTP stratum 2 servers and two switches.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 4-3 NTP Peer and Server Association



In this configuration, the switches were configured as follows:

- Stratum 2 Server 1
 - IP address -10.10.10.10
 - Stratum-2 Server-2
 - IP address -10.10.10.9
- Switch 1 IP address -10.10.10.1
- Switch 1 NTP configuration
 - NTP server 10.10.10.10
 - NTP peer 10.10.10.2
- Switch 2 IP address -10.10.10.2
- Switch 2 NTP configuration
 - NTP server 10.10.10.9
 - NTP peer 10.10.10.1

NTP Configuration Distribution

As of Cisco SAN-OS Release 2.0(1b), you can enable NTP fabric distribution for all Cisco MDS switches in the fabric. When you perform NTP configurations, and distribution is enabled, the entire server/peer configuration is distributed to all the switches in the fabric.

You automatically acquire a fabric-wide lock when you issue the first configuration command after you enabled distribution in a switch. The NTP application uses the effective and pending database model to store or commit the commands based on your configuration.

Refer to [Chapter 9, “Using the CFS Infrastructure”](#) for more information on the CFS application.

Send documentation comments to mdsfeedback-doc@cisco.com.

To enable NTP configuration fabric distribution, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# ntp distribute | Enables NTP configuration distribution to all switches in the fabric. Acquires a fabric lock and stores all future configuration changes in the pending database. |
| | switch(config)# no ntp distribute | Disables (default) NTP configuration distribution to all switches in the fabric. |

Committing NTP Configuration Changes

When you commit the NTP configuration changes, the effective database is overwritten by the configuration changes in the pending database and all the switches in the fabric receive the same configuration. When you commit the NTP configuration changes without implementing the session feature, the NTP configurations are distributed to all the switches in the fabric.

To commit the NTP configuration changes, follow these steps:

| | Command | Purpose |
|--------|-----------------------------------|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# ntp commit | Distributes the NTP configuration changes to all switches in the fabric and releases the lock. Overwrites the effective database with the changes made to the pending database. |

Discarding NTP Configuration Changes

After making the configuration changes, you can choose to discard the changes or to commit them. In either case, the lock is released.

To discard NTP configuration changes, follow these steps:

| | Command | Purpose |
|--------|----------------------------------|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# ntp abort | Discards the NTP configuration changes in the pending database and releases the fabric lock. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Releasing Fabric Session Lock

If you have performed a NTP fabric task and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.



Tip

The changes are only available in the volatile directory and are subject to being discarded if the switch is restarted.

To use administrative privileges and release a locked NTP session, use the **clear ntp session** command.

```
switch# clear ntp session
```

Database Merge Guidelines

Refer to the [“CFS Merge Support” section on page 9-7](#) for detailed concepts.

When merging two fabrics, follow these guidelines:

- Be aware that the merge is a union of the existing and the received database in each switch in the fabric.
- Do not configure an IP address as a server on one switch while and as a peer on another switch. The merge can fail if this configuration exists.
- Verify that the union of the databases does not exceed the max limit of 64.

NTP Session Status Verification

To verify the status of the NTP session, use the **show ntp session-status** command.

```
switch# show ntp session-status
last-action : Distribution Enable      Result : Success
```

Management Interface Configuration

On director class switches, a single IP address is used to manage the switch. The active supervisor module's management (mgmt0) interface uses this IP address. The mgmt0 interface on the standby supervisor module remains in an inactive state and cannot be accessed until a switchover happens. After a switchover, the mgmt0 interface on the standby supervisor module becomes active and assumes the same IP address as the previously active supervisor module.

The management interface on the switch allows multiple simultaneous Telnet or SNMP sessions. You can remotely configure the switch through the management interface, but first you must configure some IP parameters (IP address, subnet mask) so that the switch is reachable. You can manually configure the management interface from the CLI.

The management port (mgmt0) is autosensing and operates in full duplex mode at a speed of 10/100 Mbps. The speed and mode cannot be configured.

Send documentation comments to mdsfeedback-doc@cisco.com.


Note

Before you begin to configure the management interface manually, obtain the switch's IP address and IP subnet mask. Also make sure the console cable is connected to the console port.

Obtaining Remote Management Access

In some cases, a switch interface might be administratively shut down. You can check the status of an interface at any time by using the **show interface mgmt 0** command.

To obtain remote management access, follow these steps:

| | Command | Command |
|--------|---|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. You can also abbreviate the command to config t . |
| Step 2 | switch(config)# interface mgmt 0 | Enters the interface configuration mode on the specified interface (mgmt0). You can use the management Ethernet interface on the switch to configure the management interface. |
| Step 3 | switch(config)# ip address 1.1.1.0 255.255.255.0 | Enters the IP address and IP subnet mask for the interface specified in Step 2. |
| Step 4 | switch(config-if)# no shutdown | Enables the interface. |
| Step 5 | switch(config-if)# exit | Returns to configuration mode. |
| Step 6 | switch(config)# ip default-gateway 1.1.1.1 | Configures the default gateway address. |

Using the force Option

When you try to shut down a management interface (mgmt0), a follow-up message confirms your action before performing the operation. You can use the **force** option to bypass this confirmation. The following example shuts down the interface without using the **force** option:

```
switch# config t
switch(config)# interface mgmt 0
switch(config-if)# shutdown
Shutting down this interface will drop all telnet sessions.
Do you wish to continue (y/n)? y
```

The following example shuts down the interface using the **force** option:

```
switch# config t
switch(config)# interface mgmt 0
switch(config-if)# shutdown force
```


Note

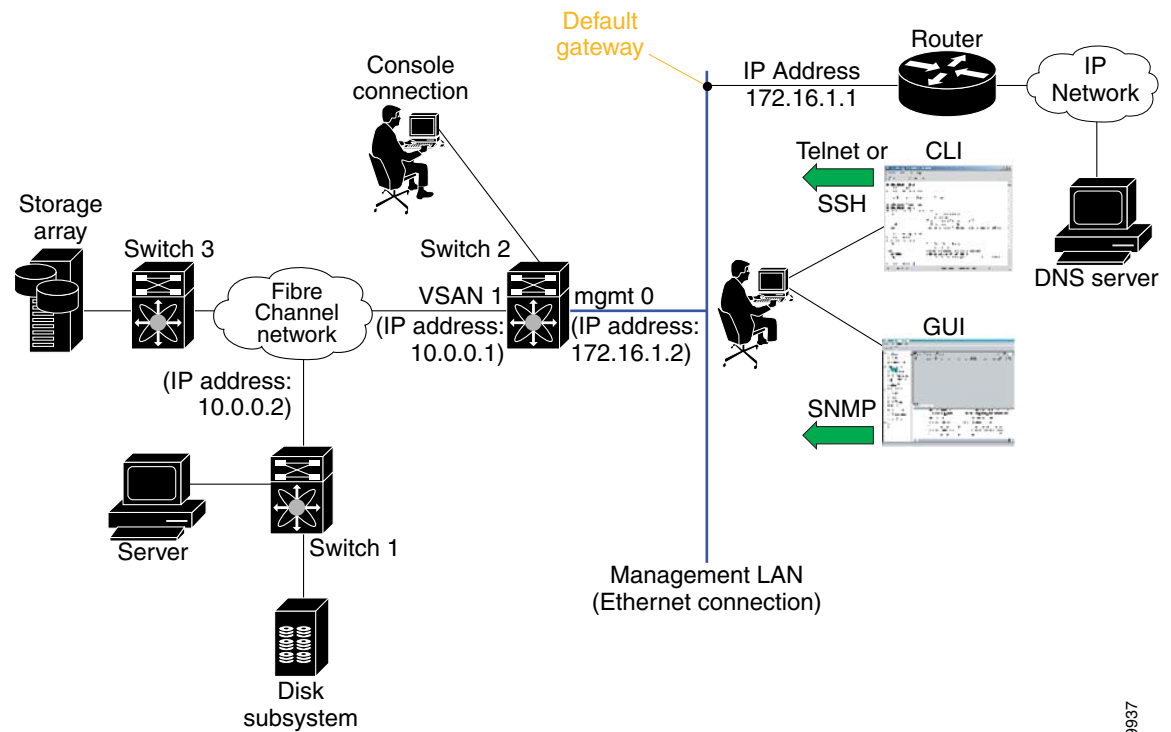
You need to explicitly configure a default gateway to connect to the switch and send IP packets or add a route for each subnet.

Send documentation comments to mdsfeedback-doc@cisco.com.

Default Gateway Configuration

The supervisor module sends IP packets with unresolved destination IP addresses to the default gateway (see Figure 4-4).

Figure 4-4 **Default Gateway**



79937

Configuring the Default Gateway

To configure the IP address of the default gateway, follow these steps:

| | Command | Purpose |
|--------|--|---------------------------------------|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# ip default-gateway 172.16.1.1 | Configures the 172.16.1.1 IP address. |

[Send documentation comments to mdsfeedback-doc@cisco.com.](mailto:mdsfeedback-doc@cisco.com)

Telnet Server Connection

The Telnet server is enabled by default on all switches in the Cisco MDS 9000 Family. If you require a secure SSH connection, you need to disable the default Telnet connection and then enable the SSH connection (see the “[Enabling SSH Service](#)” section on page 19-34).



Note

For information on connecting a terminal to the supervisor module console port, refer to the *Cisco MDS 9200 Series Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide*.



Tip

A maximum of 16 sessions are allowed in any switch in the Cisco MDS 9500 Series or the Cisco MDS 9200 Series.

Make sure the terminal is connected to the switch and that the switch and terminal are both powered on.

Disabling a Telnet Connection

To disable Telnet connections to the switch, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# no telnet server enable updated | Disables the Telnet server. |
| | switch(config)# telnet server enable updated | Enables the Telnet server to return a Telnet connection from a secure SSH connection. |

Working with Configuration Files

Configuration files can contain some or all of the commands needed to configure one or more switches. For example, you might want to download the same configuration file to several switches that have the same hardware configuration so that they have identical module and port configurations.

This section describes how to work with configuration files and has the following topics:

- [Displaying Configuration Files, page 4-25](#)
- [Downloading Configuration Files to the Switch, page 4-25](#)
- [Saving the Configuration, page 4-27](#)
- [Copying Files, page 4-28](#)
- [Backing Up the Current Configuration, page 4-29](#)
- [Rolling Back to a Previous Configuration, page 4-30](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying Configuration Files

Use the **show running-config** command to view the running configuration file.

```
switch# show running-config
Building Configuration ...
  interface port-channel 98
interface fc1/1
  interface fc1/2
interface mgmt0
ip address 172.22.95.112 255.255.255.0
no shutdown
vsan database
vsan 2
clock summer-time Pacific 1 Sun Apr 02:00 5 Sun Oct 02:00 60
switchname switch112
```

Use the **show startup-config** command to view the startup configuration file.

```
switch# show startup-config
  interface port-channel 98
  interface fc1/1
channel-group 98 force
no shutdown
  interface mgmt0
ip address 172.22.95.112 255.255.255.0
boot system system-237; ep-41
boot kickstart boot-237 ep-41
ip domain-name cisco.com
```

Downloading Configuration Files to the Switch

You can configure a switch in the Cisco MDS 9000 Family by using configuration files you create or download from another switch. In addition, you can store configuration files on a bootflash device on the supervisor module and you can configure the switch using a configuration stored on an external CompactFlash disk.

Before you begin downloading a configuration file using a remote server, do the following:

- Ensure the configuration file to be downloaded is in the correct directory on the remote server.
- Ensure that the permissions on the file are set correctly. Permissions on the file should be set to world-read.
- Ensure the switch has a route to the remote server. The switch and the remote server must be in the same subnetwork if you do not have a router or default gateway to route traffic between subnets.

Check connectivity to the remote server using the **ping** command.

Send documentation comments to mdsfeedback-doc@cisco.com.

From a Remote Server

To configure a switch in the Cisco MDS 9000 Family using a configuration file downloaded from a remote server using TFTP, FTP, SCP, or SFTP, follow these steps:

-
- Step 1** Log into the switch through the console port or through a Telnet or SSH session.
- Step 2** Configure the switch using the configuration file downloaded from the remote server using the **copy <scheme> :// <server address> system:running-config** command, where *scheme* is TFTP, FTP, SCP, or SFTP.

The configuration file downloads and the commands are executed as the file is parsed line by line.

Use the following command to download a configuration file from a remote server to the running configuration.

```
switch# copy <scheme>://<url> system:running-config
```

From an External CompactFlash Disk (slot0:)



Note

The physical media must be inserted into slot0: after you log into the switch.

To configure a switch in the Cisco MDS 9000 Family using a configuration file stored on an external CompactFlash disk, follow these steps:

-
- Step 1** Log into the switch through the console port or through a Telnet or SSH session.
- Step 2** Locate the configuration file using the **cd** and **dir** commands. (See the “Copying Files” section on [page 4-28](#).)
- Step 3** Configure the switch using the configuration file stored on the external CompactFlash disk using the **copy <source file> system:running-config** command.

The commands are executed as the file is parsed line by line.

Use the following command to download a configuration file from an external CompactFlash to the running configuration:

```
switch copy slot0:dns-config.cfg system:running-config
```

Saving Configuration Files to an External Device

You can save a configuration file stored on internal storage to a remote server or to an external Flash device on the switch.

Send documentation comments to mdsfeedback-doc@cisco.com.

To a Remote Server

To save a configuration file to a remote server such as TFTP, FTP, SCP, or SFTP, follow these steps:

-
- Step 1** Log into the switch through the console port or through a Telnet or SSH session.
 - Step 2** Save the configuration using the **copy system:running-config <scheme> :// <url>** command, where *scheme* is TFTP, FTP, SCP, or SFTP.
 - Step 3** Specify the IP address or host name of the remote server and the name of the file to download.
The configuration file is saved to the remote server.
-

Use the following command to save a running configuration file to a remote server:

```
switch# copy system:running-config <scheme>://<url>
```

Use the following command to save a startup configuration file to a remote server:

```
switch# copy nvram:startup-config <scheme>://<url>
```

To an External CompactFlash Disk (slot0:)

To save a configuration file on an external CompactFlash device, follow these steps:

-
- Step 1** Log into the switch through the console port or through a Telnet session.
 - Step 2** Locate the configuration file using the **cd** and **dir** commands. (See the “Copying Files” section on [page 4-28](#).)
 - Step 3** Save the configuration file using the **copy system:running-config <destination file>** command.
The configuration file is saved to the CompactFlash disk.
-

Use the following command to save a running configuration file to an external CompactFlash disk:

```
switch# copy system:running-config slot0:dns-config.cfg
```

Use the following command to save a startup configuration file to an external CompactFlash disk:

```
switch# copy nvram:startup-config slot0:dns-config.cfg
```

Saving the Configuration

After you have created a running configuration in system memory, you can save it to the startup configuration in NVRAM.

Use the following **copy** command to save the configuration to NVRAM:

```
switch# copy system:running-config nvram:startup-config
```

The **copy running-config startup-config** command is an alias to the previous command and is used frequently throughout this guide.

Send documentation comments to mdsfeedback-doc@cisco.com.

To cancel the copy operation initiated by another switch, use the following command:

```
switch# system startup-config abort
```

To cancel the operation locally and throughout the fabric, enter **Ctrl-c** on the console or telnet session of the initiator switch.

See the “[Preserving Module Configuration](#)” section on page 7-7.

Copying Startup Configuration to the Fabric

As of SAN-OS Release 2.1(1a), you can use Cisco Fabric Services (CFS) to instruct the other switches in the fabric to save their configurations to their local NVRAM using the following **copy** command:

```
switch# copy running-config startup-config fabric
```



Note

If any remote switch in the fabric fails to complete the **copy running-config startup-config fabric** process, the request is discarded on the initiator switch and the failure errors are displayed in the initiator switch CLI session.

Unlocking the Startup Configuration File

As of SAN-OS Release 2.0(1b), the startup configuration file can be locked by applications on the switch. To display locks on the startup configuration file, use the following command:

```
switch# show system internal sysmgr startup-config locks
```

To release a lock on the startup configuration file, use the following command:

```
switch# system startup-config unlock 10
```

Copying Files

The syntax for the **copy** command follows and is explained in [Table 4-1](#).

```
switch# copy <scheme>://<username@><server>/<file name>
<scheme>://<username@><server>/<file name>
```

Table 4-1 *copy Command Syntax*

| Scheme | Server | File Name |
|-----------|--|----------------|
| bootflash | sup-active sup-standby sup-1 or module-5 sup-2 or module-6 sup-local sup-remote | User-specified |
| slot0 | — | User-specified |
| volatile | — | User-specified |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 4-1 *copy Command Syntax (continued)*

| Scheme | Server | File Name |
|-------------------|------------------------|-----------------------------------|
| nvr | — | startup-config or snapshot-config |
| system | — | running-config |
| tftp ¹ | IP address or DNS name | User-specified |
| ftp | | |
| scp (secure copy) | | |
| sftp | | |
| core | slot-number | Process identifier number |

1. When downloading and uploading files, a TFTP limitation restricts a TFTP client to a 32 MB file size and some TFTP servers to a 16 MB file size.

- This example shows how to copy a file from the active supervisor module's (sup-1 in slot 5) bootflash to the standby supervisor module's (sup-2 in slot 6) bootflash.

```
switch# copy bootflash:system_image bootflash://sup-2/system_image
```

- This example shows how to overwrite the contents of an existing configuration in NVRAM.

```
switch# copy nvram:snapshot-config nvram:startup-config
Warning: this command is going to overwrite your current startup-config.
Do you wish to continue? {y/n} [y] y
```

- This example shows how to copy a running configuration to the bootflash: file system.

```
switch# copy system:running-config bootflash:my-config
```

- This example shows how to copy a system image file from the SCP server to bootflash.

```
switch# copy scp://user@10.1.7.2/system-image bootflash:system-image
```

- This example shows how to copy a script file from the SFTP server to the volatile: file system.

```
switch# copy sftp://172.16.10.100/myscript.txt volatile:myscript.txt
```



Note

Use the **show version image** command to verify if the downloaded images are valid.

Backing Up the Current Configuration

Before installing or migrating to any software configuration, back up the startup configuration.

- This example shows how to create a snapshot of the startup configuration in a predefined location on the switch (binary file).

```
switch# copy nvram:startup-config nvram:snapshot-config
```

- This example shows how to back up the startup configuration copy in the bootflash: file system (ASCII file).

```
switch# copy nvram:startup-config bootflash:my-config
```

Send documentation comments to mdsfeedback-doc@cisco.com.

- This example shows how to back up the startup configuration to the TFTP server (ASCII file).

```
switch# copy nvram:startup-config tftp://172.16.10.100/my-config
```
- This example shows how to back up the running configuration to the bootflash: file system (ASCII file).

```
switch# copy system:running-config bootflash:my-config
```

Rolling Back to a Previous Configuration

All switch configurations reside in the internal bootflash: file system. If your internal bootflash: file system is corrupted, you could potentially lose your configuration. Save and back up your configuration file periodically.

- This example shows how to roll back to a snapshot copy of a previously saved running configuration (binary file).

```
switch# copy nvram:snapshot-config nvram:startup-config
```



Note

You can issue a rollback command only when a snapshot is already created. Otherwise, you will receive the `No snapshot-config found` error message.

- This example shows how to roll back to a configuration copy that was previously saved in the bootflash: file system (ASCII file).

```
switch# copy bootflash:my-config nvram:startup-config
```



Note

Each time a **copy running-config startup-config** command is issued, a binary file is created and the ASCII file is updated. A valid binary configuration file reduces the overall boot time significantly. A binary file cannot be uploaded, but its contents can be used to overwrite the existing startup configuration. The **write erase** command clears the binary file.

Restoring the Configured Redundancy Mode



Tip

If you configure the **combined** mode as the redundancy mode for power supplies on a Cisco MDS 9509 switch, exert care when using the **write erase** and **reload** command sequence before rolling back to a saved configuration.

By issuing the **write erase** command and the **reload** command, you restore the switch settings to their factory defaults. This sequence also restores the redundancy mode setting for the power supplies back to the **redundant** mode (default).

Depending on the type of power supply, the input voltage, and the number of modules (line cards) in the chassis, the redundancy mode may prevent the line cards from being powered on after a system reboot (see the “[Power Supply Configuration Modes](#)” section on page 8-6). If you use this sequence, the commands that apply to the powered down line cards will not be enforced on the switch (and will not be part of its running configuration).

Send documentation comments to mdsfeedback-doc@cisco.com.

If using the **write erase** and **reload** command sequence before rolling back to a saved configuration, follow these steps:

-
- | | |
|---------------|--|
| Step 1 | Manually change (if originally configured) the redundant mode configuration to combined mode. |
| Step 2 | Wait until all modules are back online—the module status displays <code>ok</code> in response to the show module command. |
| Step 3 | Roll back to the saved configuration (see the “Rolling Back to a Previous Configuration” section on page 4-30). |
-

Downgrading from a Higher Release

Use the **install all** command to gracefully reload the switch and handle configuration conversions. When downgrading any switch in the Cisco MDS 9000 Family, avoid using the **reload** command.

For example, to revert to Cisco MDS SAN-OS Release 1.3(4b) or 1.3(5) from Release 2.x, follow these steps:

Send documentation comments to mdsfeedback-doc@cisco.com.

- Step 1** Issue the **show incompatibility system image-filename** command to determine if you need to disable any features not supported by the older release.

```
Switch# show incompatibility system bootflash:m9200-ek9-mz.1.3.4b.bin
```

The following configurations on active are incompatible with the system image

```
1) Service :port-channel , Capability :CAP_FEATURE_AUTO_CREATED_PORT_CHANNEL
   Description :active mode port channels, auto create enabled ports or auto created
   port-channels are present
   Capability requirement :STRICT
```

```
2) Service :cfs , Capability :CAP_FEATURE_CFS_ENABLED_VSAN
   Description :CFS - Distribution is enabled for VSAN
   Capability requirement :STRICT
```

```
3) Service :cfs , Capability :CAP_FEATURE_CFS_ENABLED_SYSLOGD
   Description :CFS - Distribution is enabled for SYSLOGD
   Capability requirement :STRICT
```

```
4) Service :cfs , Capability :CAP_FEATURE_CFS_ENABLED_ROLE
   Description :CFS - Distribution is enabled for ROLE
   Capability requirement :STRICT
```

```
5) Service :cfs , Capability :CAP_FEATURE_CFS_ENABLED_CALLHOME
   Description :CFS - Distribution is enabled for CALLHOME
   Capability requirement :STRICT
```

```
6) Service :cfs , Capability :CAP_FEATURE_CFS_ENABLED_PORT_SECURITY
   Description :CFS - Distribution is enabled for PORT-SECURITY
   Capability requirement :STRICT
```

```
7) Service :cfs , Capability :CAP_FEATURE_CFS_ENABLED_NTP
   Description :CFS - Distribution is enabled for NTP
   Capability requirement :STRICT
```

```
8) Service :cfs , Capability :CAP_FEATURE_CFS_ENABLED_TACACS
   Description :CFS - Distribution is enabled for TACACS
   Capability requirement :STRICT
```

```
9) Service :cfs , Capability :CAP_FEATURE_CFS_ENABLED_RADIUS
   Description :CFS - Distribution is enabled for RADIUS
   Capability requirement :STRICT
```

```
10) Service :cfs , Capability :CAP_FEATURE_CFS_ENABLED_SFM
   Description :CFS - Distribution is enabled for SFM
   Capability requirement :STRICT
```

```
11) Service :cfs , Capability :CAP_FEATURE_CFS_ENABLED_DEVICE_ALIAS
   Description :CFS - Distribution is enabled for DEVICE-ALIAS
   Capability requirement :STRICT
```

- Step 2** Save the configuration using the **copy running-config startup-config** command.

- Step 3** Issue the **install all** command to to downgrade the software (see the [“Automated Upgrades” section on page 6-5](#)).

Send documentation comments to mdsfeedback-doc@cisco.com.

Accessing Remote File Systems

To access contents of the standby supervisor module (remote), follow these steps:

Step 1 Verify if the standby supervisor module has sufficient space for new image files.

```
switch# dir bootflash://sup-remote
12198912    Aug 27 17:21:10 2003  bt-39a
12198912    Aug 27 16:29:18 2003  m9500-sf1ek9-kickstart-mzg.1.3.0.39a.bin
1921922    Sep 14 19:58:12 2003  aOldImage
1864931    Apr 29 12:41:50 2003  bOldImage
1864931    Apr 29 12:41:59 2003  dplug2
12288      Apr 18 20:23:11 2003  lost+found/
12097024    Nov 21 16:34:18 2003  m9500-sf1ek9-kickstart-mz.1.3.1.1.bin
41574014    Nov 21 16:34:47 2003  m9500-sf1ek9-mz.1.3.1.1.bin
1024       Oct 28 20:24:59 2003  newer-fs/
2021518    Oct 11 15:49:41 2003  plugin-69a
Usage for bootflash://sup-remote
102081536 bytes used
82478080 bytes free
184559616 bytes total
```

Step 2 Delete files, if required, to make more space for the new image files.

```
switch# del aOldImage
```

Deleting Files

Assuming you are already in the bootflash: file system, use the **delete** command as follows:

- This example shows how to delete a file from the bootflash: file system.

```
switch# delete dns_config.cfg
```

- This example shows how to delete a file from an external CompactFlash (slot0:) file system.

```
switch# delete slot0:dns_config.cfg
```

- This example shows how to delete the file named test from the Flash card inserted in slot 0.

```
switch# delete slot0:test
Delete slot0:test? [y/n]: y
```

- This example shows how to delete the entire my-dir directory and all its contents.

```
switch# delete bootflash:my-dir
```

[Send documentation comments to mdsfeedback-doc@cisco.com.](mailto:mdsfeedback-doc@cisco.com)

Configuring Console Port Settings

A console port is an asynchronous serial port that enables switches in the Cisco MDS 9000 Family to be set up for initial configuration through a standard RS-232 port with an RJ-45 connector. Any device connected to this port must be capable of asynchronous transmission. Connection to a terminal requires a terminal emulator to be configured as 9600 baud, 8 data bits, 1 stop bit, no parity.



Caution

The console baud rate automatically reverts to the default rate (9600) after any BIOS upgrade.

To configure the console port parameters from the console terminal, follow these steps:

| | Command | Command |
|--------|--|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# line console switch(config-console)# | Enters the line console configuration mode. |
| Step 3 | switch(config-console)# speed 9600 | Configures the port speed for the serial console. The default console baud rate is 9600 baud. The valid range is between 110 and 115,200 bps (110, 150, 300, 600, 1200, 2400, 4800, 9600, 19200, 28800, 38400, 57600, 115200). Be sure to specify one of these exact values. |
| Step 4 | switch(config-console)# databits 8 | Configures the data bits for the console connection. The default is 8 data bits and the valid range is between 5 and 8 data bits. |
| Step 5 | switch(config-console)# stopbits 1 | Configures the stop bits for the console connection. The default is 1 stop bit and the valid values are 1 or 2 stop bits. |
| Step 6 | switch(config-console)# parity none | Configures the parity for the console connection. The default is no parity and the valid values are even or odd parity. |

Verifying Console Port Settings

Use the **show line console** command to verify the configured console settings. This command also displays problems that may have occurred along with the other registration statistics.

```
switch# show line console
line Console:
  Speed:          9600 bauds
  Databits:       8 bits per byte
  Stopbits:       1 bit(s)
  Parity:         none
  Modem In: Enable
  Modem Init-String -
    default : ATE0Q1&D2&C1S0=1\015
  Statistics: tx:12842   rx:366   Register Bits:RTS|CTS|DTR|DSR|CD|RI
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Configuring COM1 Port Settings

A COM1 port is a RS-232 port with a DB-9 interface that enables you to connect to an external serial communication device such as a modem. Connection to a terminal requires the terminal emulator to be configured as 9600 baud, 8 data bits, 1 stop bit, no parity.

To configure the COM1 port settings, follow these steps:

| | Command | Command |
|--------|--|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# line com1 switch(config-com1)# | Enters the COM1 port configuration mode. |
| Step 3 | switch(config-com1)# speed 9600 | Configures the port speed for the COM1 connection. The default console baud rate is 9600 baud. The valid range is between 110 and 115,200 bps (110, 150, 300, 600, 1200, 2400, 4800, 9600, 19200, 28800, 38400, 57600, 115200). Be sure to specify one of these exact values. Note This configuration depends on the incoming speed of the modem connected to COM1. |
| Step 4 | switch(config-com1)# databits 8 | Configures the data bits for the COM1 connection. The default is 8 data bits and the valid range is between 5 and 8 data bits. |
| Step 5 | switch(config-com1)# stopbits 1 | Configures the stop bits for the COM1 connection. The default is 1 stop bits and the valid values are 1 or 2 stop bits. |
| Step 6 | switch(config-com1)# parity none | Configures the parity for the COM1 connection. The default is no parity and the valid values are even or odd parity. |
| Step 7 | switch(config-com1)# no flowcontrol hardware | Disables hardware flow control. By default, hardware flow control is enabled on all switches in the Cisco 9000 Family. When enabled, this option is useful in protecting data loss at higher baud rates. Note This option is only available through the COM1 port. |

Verifying COM1 Port Settings

Use the **show line com1** command to verify the configured COM1 settings. This command also displays problems that may have occurred along with the other registration statistics.

```
switch# show line com1
line Aux:
  Speed:          9600 bauds
  Databits:       8 bits per byte
  Stopbits:       1 bit(s)
  Parity:         none
  Modem In: Enable
  Modem Init-String -
    default : ATE0Q1&D2&C1S0=1\015
  Statistics: tx:17    rx:0    Register Bits:RTS|DTR
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Configuring Modem Connections

Modems can only be configured if you are connected to the console or COM1 ports. A modem connection to a switch in the Cisco MDS 9000 Family does not affect switch functionality.



Note

If you plan on connecting a modem to the console port or the COM1 port of a switch in the Cisco MDS 9000 Family, refer to the *Cisco MDS 9200 Series Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide*. COM1 ports are not available on switches in the Cisco MDS 9100 Series. Refer to the *Cisco MDS 9100 Series Hardware Installation Guide*.

Guidelines to Configure Modems



Tip

We recommend you use the COM1 port to connect the modem from any director in the Cisco MDS 9500 Series or any Switch in the Cisco MDS 9200 Series.

The following guidelines apply to modem configurations:

- The following Cisco modems were tested to work in the Cisco SAN-OS environment:
 - MultiTech MT2834BA (<http://www.multitech.com/PRODUCTS/Families/MultiModemII/>)
 - Hayes Accura V.92 (<http://www.hayesmicro.com/Products/accura-prod-v92.htm>)
- Connect the modem before attempting to configure the modem.
- Do not connect a modem to the console port while the system is booting.

Follow the procedure specified in the “[Initializing a Modem in a Powered-On Switch](#)” section on [page 4-39](#).

Enabling Modem Connections

To configure a modem connection through the COM1 port, follow these steps:

| | Command | Command |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# line com1 switch(config-com1)# | Enters the COM1 port configuration mode. |
| Step 3 | switch(config-com1)# modem in | Enables the COM1 port to only connect to a modem. |
| | switch(config-com1)# no modem in | Disables (default) the current modem from executing its functions. |

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure a modem connection through the console port, follow these steps:

| | Command | Command |
|--------|--|--|
| Step 1 | switch# confi g t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# line console switch(config-console)# | Enters the console port configuration mode. |
| Step 3 | switch(config-console)# modem in | Enables the console port to only connect to a modem. |
| | switch(config-console)# no modem in | Disables (default) the current modem from executing its functions. |

Configuring the Initialization String

Switches in the Cisco MDS 9500 Series and the Cisco MDS 9200 Series have a default initialization string (ATE0Q1&D2&C1S0=1\015) to detect connected modems. The default string detects connected modems supported by Cisco Systems. The default string contents are as follows:

- AT—Attention
- E0 (required)—No echo
- Q1—Result code on
- &D2—Normal data terminal ready (DTR) option
- &C1—Enable tracking the state of the data carrier.
- S0=1—Pick up after one ring
- \015 (required)—carriage return in octal

You may retain the default string or change it to another string (80 character limit) using the **user-input** option. This option is provided if you prefer to use a modem that is not supported or tested by Cisco systems. If you change the string, the changes you make are permanent and remain in effect unless you change them again. Rebooting the system or restarting the CLI does not change the modem initialization string. The switch is not affected even if the modem is not functioning.



Tip

We recommend you use the default initialization string. If the required options are not provided in the user-input string, the initialization string is not processed.

The modem initialization string usage depends on the modem state when the switch boots:

- If the modem is already attached to the switch during boot-up, the default initialization string is written to the modem (see the [“Configuring the Default Initialization String”](#) section on page 4-38).
- If the modem is not attached to the switch during boot-up, then attach the modem as outlined in the *Cisco MDS 9000 Family Hardware Installation Guide* (depending on the product), and follow the procedure provided in this section (see the [“Configuring a User-Specified Initialization String”](#) section on page 4-38).



Note

You can perform the configuration specified in this section only if you are connected to the console port or the COM1 port.

Send documentation comments to mdsfeedback-doc@cisco.com.

Configuring the Default Initialization String

To configure the default initialization string through the COM1 port, follow these steps:

| | Command | Command |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# line com1 switch(config-com1)# | Enters the COM1 port configuration mode. |
| Step 3 | switch(config-com1)# modem init-string default | Writes the default initialization string to the modem. |

To configure the default initialization string through the console port, follow these steps:

| | Command | Command |
|--------|---|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# line com1 switch(config-console)# | Enters the console port configuration mode. |
| Step 3 | switch(config-console)# modem init-string default | Writes the default initialization string to the modem. |

Configuring a User-Specified Initialization String

To configure a user-specified initialization string through the COM1 port, follow these steps:

| | Command | Command |
|--------|--|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# line com1 switch(config-com1)# | Enters the COM1 port configuration mode. |
| Step 3 | switch(config-com1)# modem set-string user-input ATE0Q1&D2&C1S0=3\015 | Assigns the user-specified initialization string to its corresponding profile. |
| | switch(config-com1)# no modem set-string | Note You must first set the user-input string, before initializing the string. Reverts the configured initialization string to the factory default string. |
| Step 4 | switch(config-com1)# modem init-string user-input | Writes the user-specified initialization string to the modem. |

To configure a user-specified initialization string through the console port, follow these steps:

| | Command | Command |
|--------|---|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# line com1 switch(config-console)# | Enters the console port configuration mode. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Command |
|--------|---|---|
| Step 3 | switch(config-console)# modem set-string user-input ATE0Q1&D2&C1S0=3\015 | Assigns the user-specified initialization string to its corresponding profile. |
| | | Note You must first set the user-input string, before initializing the string. |
| | switch(config-com1)# no modem set-string | Reverts the configured initialization string to the factory default string. |
| Step 4 | switch(config-console)# modem init-string user-input | Writes the user-specified initialization string to the modem. |

Initializing a Modem in a Powered-On Switch

When a switch is already powered-on and the modem is later connected to either the console port or the COM1 port, you can initialize the modem using the **modem connect line** command in EXEC mode. You can specify the **com1** option if the modem is connected to the COM1 port, or the **console** option if the modem is connected to the console.

To connect a modem to a switch that is already powered on, follow these steps.

-
- Step 1** Wait until the system has completed the boot sequence and the system image is running.
 - Step 2** Connect the modem to the switch as specified in the *Cisco MDS 9200 Series Hardware Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide*.
 - Step 3** Initialize the modem using the **modem connect line** command in EXEC mode.
-

Verifying the Modem Connection Configuration

Use the **show line** command to verify the configured modem settings.

```
switch# show line
line Console:
  Speed:          9600 bauds
  Databits:       8 bits per byte
  Stopbits:      1 bit(s)
  Parity:         none
  Modem In: Enable
  Modem Init-String -
    default : ATE0Q1&D2&C1S0=1\015
  Statistics: tx:12842    rx:366    Register Bits:RTS|CTS|DTR|DSR|CD|RI
line Aux:
  Speed:          9600 bauds
  Databits:       8 bits per byte
  Stopbits:      1 bit(s)
  Parity:         none
  Modem In: Enable
  Modem Init-String -
    default : ATE0Q1&D2&C1S0=1\015
  Statistics: tx:17      rx:0      Register Bits:RTS|DTR
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Configuring CDP

The Cisco Discovery Protocol (CDP) is an advertisement protocol used by Cisco devices to advertise itself to other Cisco devices in the same network. CDP runs on the data link layer and is independent of Layer 3 protocols. Cisco devices that receive the CDP packets cache the information to make it is accessible through the CLI and SNMP.

CDP is supported on the management Ethernet interface on the supervisor module and the Gigabit Ethernet interface on the IPS module. The CDP daemon is restartable and switchable. The running and startup configurations are available across restarts and switchovers.

CDP version 1 (v1) and version 2 (v2) are supported in Cisco MDS 9000 Family switches. CDP packets with any other version number are silently discarded when received.

When the interface link is established, CDP is enabled by default and three CDP packets are sent at one-second intervals. Following this, the CDP frames are sent at the globally-configured refresh interval.

To globally disable the CDP, follow these steps:

| | Command | Command |
|--------|--|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# no cdp enable Operation in progress. Please check global parameters switch(config-console)# | Disables the CDP protocol on the switch. When CDP is disabled on an interface, one packet is sent to clear out the switch state with each of the receiving devices. |
| | switch(config)# cdp enable Operation in progress. Please check global parameters switch(config)# | Enables (default) the CDP protocol on the switch. When CDP is enabled on an interface, one packet is sent immediately. Subsequent packets are sent at the configured refresh time. |

To disable the CDP protocol on a specific interface, follow these steps:

| | Command | Command |
|--------|--|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# interface gigabitethernet 8/8 switch(config-if)# | Configures the Gigabit Ethernet interface for the module in slot 8 port 8. |
| Step 3 | switch(config-if)# no cdp enable Operation in progress. Please check interface parameters switch(config-console)# | Disables the CDP protocol on the selected interface. When CDP is disabled on an interface, one packet is sent to clear out the switch state with each of the receiving devices. |
| | switch(config-if)# cdp enable Operation in progress. Please check interface parameters switch(config)# | Enables (default) the CDP protocol on the selected interface. When CDP is enabled on an interface, one packet is sent immediately. Subsequent packets are sent at the configured refresh time. |

Send documentation comments to mdsfeedback-doc@cisco.com.

To globally configure the refresh time interval for the CDP protocol, follow these steps:

| | Command | Command |
|--------|--|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# cdp timer 100 switch(config)# | Sets the refresh time interval in seconds. The default is 60 seconds and the valid range is from 5 to 255 seconds. |
| | switch(config)# no cdp timer 100 switch(config)# | Reverts the refresh time interval to the factory default of 60 seconds. |

To globally configure the hold time advertised in CDP packets, follow these steps:

| | Command | Command |
|--------|---|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# cdp holdtime 200 switch(config)# | Sets the hold time advertised in CDP packets in seconds. The default is 180 seconds and the valid range is from 10 to 255 seconds. |
| | switch(config)# no cdp holdtime 200 switch(config)# | Reverts the hold time to the factory default of 180 seconds. |

To globally configure the CDP version, follow these steps:

| | Command | Command |
|--------|--|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# cdp advertise v1 switch(config)# | Sets the CDP version to be used. The default is version 2 (v2). The valid options are v1 and v2. |
| | switch(config)# no advertise v1 switch(config)# | Reverts the version to the factory default of v2. |

Clearing CDP Counters and Tables

Use the **clear cdp counters** command to clear CDP traffic counters for all interfaces. You can issue this command for a specified interface or for all interfaces (management and Gigabit Ethernet interfaces).

```
switch# clear cdp counters
switch#
```

Use the **clear cdp table** command to clear neighboring CDP entries for all interfaces. You can issue this command for a specified interface or for all interfaces (management and Gigabit Ethernet interfaces).

```
switch# clear cdp table interface gigabitethernet 4/1
switch#
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying CDP Information

Use the **show cdp** command to display CDP entries. See Examples 4-1 to 4-11.

Example 4-1 *Displays All CDP Capable Interfaces and Parameters*

```
switch# show cdp all
GigabitEthernet4/1 is up
  CDP enabled on interface
  Sending CDP packets every 60 seconds
  Holdtime is 180 seconds
GigabitEthernet4/8 is down
  CDP enabled on interface
  Sending CDP packets every 60 seconds
  Holdtime is 180 seconds
mgmt0 is up
  CDP enabled on interface
  Sending CDP packets every 100 seconds
  Holdtime is 200 seconds
```

Example 4-2 *Displays All CDP Neighbor Entries*

```
switch# show cdp entry all
-----
Device ID:069038747(Kiowa3)
Entry address(es):
  IP Address: 172.22.92.5
Platform: WS-C5500, Capabilities: Trans-Bridge Switch
Interface: mgmt0, Port ID (outgoing port): 5/22
Holdtime: 136 sec

Version:
WS-C5500 Software, Version McpSW: 2.4(3) NmpSW: 2.4(3)
Copyright (c) 1995-1997 by Cisco Systems

Advertisement Version: 1
```

Example 4-3 *Displays the Specified CDP Neighbor*

```
switch# show cdp entry name 0
-----
Device ID:0
Entry address(es):
  IP Address: 0.0.0.0
Platform: DS-X9530-SF1-K9, Capabilities: Host
Interface: GigabitEthernet4/1, Port ID (outgoing port): GigabitEthernet4/1
Holdtime: 144 sec

Version:
1.1(0.144)

Advertisement Version: 2
Duplex: full
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 4-4 Displays Global CDP Parameters

```
switch# show cdp global
Global CDP information:
  CDP enabled globally
  Sending CDP packets every 60 seconds
  Sending a holdtime value of 180 seconds
  Sending CDPv2 advertisements is enabled
```

Example 4-5 Displays CDP Parameters for the Management Interface

```
switch# show cdp interface mgmt 0
mgmt0 is up
  CDP enabled on interface
  Sending CDP packets every 60 seconds
  Holdtime is 180 seconds
```

Example 4-6 Displays CDP Parameters for the Gigabit Ethernet Interface

```
switch# show cdp interface gigabitethernet 4/1
GigabitEthernet4/1 is up
  CDP enabled on interface
  Sending CDP packets every 80 seconds
  Holdtime is 200 seconds
```

Example 4-7 Displays CDP Neighbors (in brief)

```
switch# show cdp neighbors
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater
```

| Device ID | Local Intrfce | Hldtme | Capability | Platform | Port ID |
|-------------------|---------------|--------|------------|---------------|---------|
| 0 | Gig4/1 | 135 | H | DS-X9530-SF1- | Gig4/1 |
| 069038732 (Kiowa2 | mgmt0 | 132 | T S | WS-C5500 | 8/11 |
| 069038747 (Kiowa3 | mgmt0 | 156 | T S | WS-C5500 | 6/20 |
| 069038747 (Kiowa3 | mgmt0 | 158 | T S | WS-C5500 | 5/22 |

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 4-8 Displays CDP Neighbors (in detail)

```
switch# show CDP neighbor detail
-----
Device ID:0
Entry address(es):
  IP Address: 0.0.0.0
Platform: DS-X9530-SF1-K9, Capabilities: Host
Interface: GigabitEthernet4/1, Port ID (outgoing port): GigabitEthernet4/1
Holdtime: 162 sec

Version:
1.1(0.144)

Advertisement Version: 2
Duplex: full
-----
Device ID:069038732(Kiowa2)
Entry address(es):
  IP Address: 172.22.91.5
Platform: WS-C5500, Capabilities: Trans-Bridge Switch
Interface: mgmt0, Port ID (outgoing port): 8/11
Holdtime: 132 sec

Version:
WS-C5500 Software, Version MpsSW: 2.4(3) NmpSW: 2.4(3)
Copyright (c) 1995-1997 by Cisco Systems
Advertisement Version: 1
```

Example 4-9 Displays the Specified CDP Neighbor (in detail)

```
switch# show CDP neighbors interface gigabitethernet 4/1 detail
-----
Device ID:0
Entry address(es):
  IP Address: 0.0.0.0
Platform: DS-X9530-SF1-K9, Capabilities: Host
Interface: GigabitEthernet4/1, Port ID (outgoing port): GigabitEthernet4/1
Holdtime: 144 sec

Version:
1.1(0.144)

Advertisement Version: 2
Duplex: full
```


Send documentation comments to mdsfeedback-doc@cisco.com.

Example 4-10 Displays CDP Traffic Statistics for the Management Interface

```
switch# show cdp traffic interface mgmt 0
-----
Traffic statistics for mgmt0
Input Statistics:
  Total Packets: 1148
  Valid CDP Packets: 1148
    CDP v1 Packets: 1148
    CDP v2 Packets: 0
  Invalid CDP Packets: 0
    Unsupported Version: 0
    Checksum Errors: 0
    Malformed Packets: 0
Output Statistics:
  Total Packets: 2329
    CDP v1 Packets: 1164
    CDP v2 Packets: 1165
  Send Errors: 0
```

Example 4-11 Displays CDP Traffic Statistics for the Gigabit Ethernet Interface

```
switch# show cdp traffic interface gigabitethernet 4/1
-----
Traffic statistics for GigabitEthernet4/1
Input Statistics:
  Total Packets: 674
  Valid CDP Packets: 674
    CDP v1 Packets: 0
    CDP v2 Packets: 674
  Invalid CDP Packets: 0
    Unsupported Version: 0
    Checksum Errors: 0
    Malformed Packets: 0
Output Statistics:
  Total Packets: 674
    CDP v1 Packets: 0
    CDP v2 Packets: 674
  Send Errors: 0
```

Send documentation comments to mdsfeedback-doc@cisco.com.



Configuring High Availability

The Cisco MDS 9500 Series of multilayer directors support application restartability and nondisruptive supervisor switchability. The switches are protected from system failure by redundant hardware components and a high availability software framework.

This chapter includes the following sections:

- [About High Availability, page 5-1](#)
- [Switchover Mechanisms, page 5-2](#)
- [Switchover Guidelines, page 5-3](#)
- [Process Restartability, page 5-4](#)
- [Synchronizing Supervisor Modules, page 5-4](#)
- [Copying Boot Variable Images to the Standby Supervisor, page 5-4](#)
- [Displaying HA Information, page 5-5](#)

About High Availability

The high availability (HA) software framework provides the following:

- Ensures nondisruptive software upgrade capability. See [Chapter 6, “Software Images.”](#)
- Provides redundancy for supervisor module failure by using dual supervisor modules.
- Performs nondisruptive restarts of a failed process on the same supervisor module. A service running on the supervisor modules and on the switching module tracks the HA policy defined in the configuration and takes action based on this policy. This feature is also available in switches in the Cisco MDS 9100 Series and the Cisco MDS 9200 Series.
- Protects against link failure using the PortChannel (port aggregation) feature. This feature is also available in switches in the Cisco MDS 9200 Series and in the Cisco MDS 9100 Series. See [Chapter 14, “Configuring PortChannels.”](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

- Provides management redundancy using the Virtual Router Redundancy Protocol (VRRP). This feature is also available in switches in the Cisco MDS 9100 Series and in the Cisco MDS 9200 Series.

See the “[The Virtual Router Redundancy Protocol](#)” section on page 26-19.

- Provides switchovers if the active supervisor fails, the standby supervisor, if present, takes over without disrupting storage or host traffic.

Directors in the Cisco MDS 9500 Series have two supervisor modules in the two center slots (sup-1 and sup-2). When the switch powers up and both supervisor modules are present, the supervisor module that comes up first enters the active mode and the supervisor module that comes up second enters the standby mode. If both supervisor modules come up at the same time, sup-1 becomes active. The standby module constantly monitors the active module. If the active module fails, the standby module takes over without any impact to user traffic.

Switchover Mechanisms

If the active supervisor module fails, the standby module automatically takes over. You can manually initiate a switchover from an active supervisor module to a standby supervisor module.

Once a switchover process has started another switchover process cannot be started on the same switch until a stable standby supervisor module is available.



Caution

If the supervisor modules are not in a stable state (online or powered down), a switchover will not be performed.

HA Switchover Characteristics

An HA switchover has the following characteristics:

- It is stateful (nondisruptive) because control traffic is not impacted.
- It does not impact data traffic because the switching modules are not impacted.
- Switching modules are not reset.

Initiating a Switchover

To manually initiate a switchover from an active supervisor module to a standby supervisor module, issue the **system switchover** command. Once issued, another switchover process cannot be started on the same switch until a stable standby module is available.

To ensure that an HA switchover is possible, issue the **show system redundancy status** command or the **show module** command. If the command output displays the `HA-standby` state for the standby supervisor module, then the switchover is possible.

Send documentation comments to mdsfeedback-doc@cisco.com.

Switchover Guidelines

Be aware of the following guidelines when performing a switchover:

- When you manually initiate a switchover, system messages indicate the presence of two supervisor modules.
- A switchover can only be performed when two supervisor modules are functioning in the switch.
- The modules in the chassis are functioning as designed.

Verifying Switchover Possibilities

This section describes how to verify the status of the switch and the modules before a switchover.

- Use the **show system redundancy status** command to ensure that the system is ready to accept a switchover.
- Use the **show module** command to verify the status (and presence) of a module at any time. A sample output of the **show module** command follows:

```
switch# show module
Mod  Ports  Module-Type                Model                Status
---  ---
2    8      IP Storage Services Module DS-X9308-SMIP        ok
5    0      Supervisor/Fabric-1       DS-X9530-SF1-K9      active *
6    0      Supervisor/Fabric-1       DS-X9530-SF1-K9      ha-standby
8    0      Caching Services Module   DS-X9560-SMAP        ok
9    32     1/2 Gbps FC Module        DS-X9032              ok

Mod  Sw          Hw          World-Wide-Name(s) (WWN)
---  ---
2    1.3(0.106a) 0.206       20:41:00:05:30:00:00:00 to 20:48:00:05:30:00:00:00
5    1.3(0.106a) 0.602       --
6    1.3(0.106a) 0.602       --
8    1.3(0.106a) 0.702       --
9    1.3(0.106a) 0.3         22:01:00:05:30:00:00:00 to 22:20:00:05:30:00:00:00

Mod  MAC-Address(es)                Serial-Num
---  ---
2    00-05-30-00-9d-d2 to 00-05-30-00-9d-de JAB064605a2
5    00-05-30-00-64-be to 00-05-30-00-64-c2 JAB06350B1R
6    00-d0-97-38-b3-f9 to 00-d0-97-38-b3-fd JAB06350B1R
8    00-05-30-01-37-7a to 00-05-30-01-37-fe JAB072705ja
9    00-05-30-00-2d-e2 to 00-05-30-00-2d-e6 JAB06280ae9
```

* this terminal session

The **Status** column in the output should display an **OK** status for switching modules and an **active** or **HA-standby** status for supervisor modules. If the status is either **OK** or **active**, you can continue with your configuration.

Send documentation comments to mdsfeedback-doc@cisco.com.

Process Restartability

Process restartability provides the high availability functionality in Cisco MDS 9000 Family switches. It ensures that process-level failures do not cause system-level failures. It also restarts the failed processes automatically. This vital process functions on infrastructure that is internal to the switch.

See the “[Displaying System Processes](#)” section on page 41-1.

Synchronizing Supervisor Modules

The running image is automatically synchronized in the standby supervisor module by the active supervisor module. The boot variables are synchronized during this process.

The standby supervisor module automatically synchronizes its image with the running image on the active supervisor module.

See the “[Replacing Modules](#)” section on page 6-24.

Copying Boot Variable Images to the Standby Supervisor

You can copy the boot variable images that are in the active supervisor module (but not in the standby supervisor module) to the standby supervisor module. Only those KICKSTART and SYSTEM boot variables that are set for the standby supervisor module can be copied. For module (line card) images, all boot variables are copied to the corresponding standby locations (bootflash or slot0) if not already present.

Automatic Copying of Boot Variables

To enable or disable automatic copying of boot variables, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# boot auto-copy | Enables automatic copying of boot variables from the active supervisor module to the standby supervisor module. |
| | switch(config)# no boot auto-copy | Disables the automatic copy feature (default). |

Verifying the Copied Boot Variables

Use the **show boot auto-copy** command to verify the current state of the copied boot variables (see [Example 5-1](#) and [Example 5-2](#)).

Example 5-1 Displays the auto-copy Option in an Enabled State

```
switch# show boot auto-copy
Boot variables Auto-Copy ON
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 5-2 *Displays the auto-copy Option in a Disabled State*

```
switch# show boot auto-copy
Boot variables Auto-Copy OFF
```

Use the **show boot auto-copy list** command to verify what files are being copied. [Example 5-3](#) displays the image being copied to the standby supervisor module's bootflash. Once this is successful, the next file will be image2.bin. This command only displays files on the active supervisor module.

Example 5-3 *Displays the Files Being Copied*

```
switch# show boot auto-copy list
File: /bootflash:/image1.bin
Bootvar: kickstart

File:/bootflash:/image2.bin
Bootvar: system
```

[Example 5-4](#) displays a typical message when the **auto-copy** option is disabled or if no files are copied.

Example 5-4 *Displays the Current auto-copy State*

```
switch# show boot auto-copy list
No file currently being auto-copied
```

Displaying HA Information

Use the **show system redundancy status** command to view the high availability status of the system (see [Example 5-5](#)). Tables [5-1](#) to [5-3](#) explain the possible output values for the redundancy, supervisor, and internal states.

Example 5-5 *Displays Redundancy Status*

```
switch# show system redundancy status
Redundancy mode
-----
      administrative:  HA
      operational:    HA
This supervisor (sup-1)
-----
      Redundancy state: Active
      Supervisor state: Active
      Internal state:   Active with HA standby
Other supervisor (sup-2)
-----
      Redundancy state: Standby
      Supervisor state: HA standby
      Internal state:   HA standby
```

The following conditions identify when automatic synchronization is possible:

- If the internal state of one supervisor module is **Active with HA standby** and the other supervisor module is **HA-standby**, the switch is operationally **HA** and can do automatic synchronization.
- If the internal state of one of the supervisor modules is **none**, the switch cannot do automatic synchronization.

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 5-1 lists the possible values for the redundancy states.

Table 5-1 Redundancy States

| State | Description |
|--------------|--|
| Not present | The supervisor module is not present or is not plugged into the chassis. |
| Initializing | The diagnostics have passed and the configuration is being downloaded. |
| Active | The active supervisor module and the switch is ready to be configured. |
| Standby | A switchover is possible. |
| Failed | The switch detects a supervisor module failure on initialization and automatically attempts to power-cycle the module three (3) times. After the third attempt it continues to display a failed state. |
| Offline | The supervisor module is intentionally shut down for debugging purposes. |
| At BIOS | The switch has established connection with the supervisor and the supervisor module is performing diagnostics. |
| Unknown | The switch is in an invalid state. If it persists, call TAC. |

Table 5-2 lists the possible values for the supervisor module state.

Table 5-2 Supervisor States

| State | Description |
|------------|---|
| Active | The active supervisor module in the switch is ready to be configured. |
| HA standby | A switchover is possible. |
| Offline | The switch is intentionally shut down for debugging purposes. |
| Unknown | The switch is in an invalid state and requires a support call to TAC. |

Table 5-3 lists the possible values for the internal redundancy states.

Table 5-3 Internal States

| State | Description |
|--------------------------------|--|
| HA standby | The HA switchover mechanism in the standby supervisor module is enabled (see the “HA Switchover Characteristics” section on page 5-2). |
| Active with no standby | A switchover is possible. |
| Active with HA standby | The active supervisor module in the switch is ready to be configured. The standby module is in the HA-standby state. |
| Shutting down | The switch is being shut down. |
| HA switchover in progress | The switch is in the process of changing over to the HA switchover mechanism. |
| Offline | The switch is intentionally shut down for debugging purposes. |
| HA synchronization in progress | The standby supervisor module is in the process of synchronizing its state with the active supervisor modules. |
| Standby (failed) | The standby supervisor module is not functioning. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 5-3 ***Internal States (continued)***

| State | Description |
|----------------------------|--|
| Active with failed standby | The active supervisor module and the second supervisor module is present but is not functioning. |
| Other | The switch is in a transient state. If it persists, call TAC. |

Send documentation comments to mdsfeedback-doc@cisco.com.



Software Images

This chapter describes how to install and upgrade software images. It includes the following sections:

- [About Software Images, page 6-1](#)
- [Essential Upgrade Prerequisites, page 6-2](#)
- [Software Upgrade Methods, page 6-4](#)
- [Automated Upgrades, page 6-5](#)
- [Upgrade Status Verification, page 6-17](#)
- [Manual Upgrade on a Dual Supervisor Switch, page 6-18](#)
- [Quick Upgrade, page 6-23](#)
- [Maintaining Supervisor Modules, page 6-23](#)
- [Replacing Modules, page 6-24](#)
- [Corrupted Bootflash Recovery, page 6-25](#)
- [Default Settings, page 6-34](#)

About Software Images

Each switch is shipped with a Cisco MDS SAN-OS operating system for Cisco MDS 9000 Family switches. The Cisco SAN-OS consists of two images—the kickstart image and the system image. To upgrade the switch to a new image, you must specify the variables that direct the switch to the images.

- To select the kickstart image, use the KICKSTART variable.
- To select the system image, use the SYSTEM variable.

The images and variables are important factors in any install procedure. You must specify the variable and the image to upgrade your switch. Both images are not always required for each install.



Note

Unless explicitly stated, the software install procedures in this section apply to any switch in the Cisco MDS 9000 Family.

Send documentation comments to mdsfeedback-doc@cisco.com.

Dependent Factors for Software Installation

The software image install procedure is dependent on the following factors:

- Software images—The kickstart and system image files reside in directories or folders that can be accessed from the Cisco MDS 9000 Family switch prompt.
- Image version—Each image file has a version.
- Flash disks on the switch—The bootflash: resides on the supervisor and the CompactFlash disk is inserted into the slot0: device.
- Supervisor modules—There are single or dual supervisor modules.

Essential Upgrade Prerequisites

Before attempting to migrate to any software image version, follow these guidelines:

- Customer Service

Before performing any software upgrade, contact your respective customer service representative to review your software upgrade requirements and to provide recommendations based on your current operating environment.



Note

If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco Systems, contact Cisco Technical Support at this URL: <http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

- Scheduling

Schedule the upgrade when the fabric is stable and steady. Ensure that everyone who has access to the switch or the network is not configuring the switch or the network during this time. All configurations are disallowed at this time.

- Space

Verify that sufficient space is available in the location where you are copying the images. This location includes the active and standby supervisor bootflash: (internal to the switch).

- Standby supervisor module bootflash: filesystem (see the [Chapter 4, “Initial Configuration”](#)).
- Internal bootflash offers approximately 200 MB of user space.

- Hardware

Avoid power interruption during any install procedure. These kinds of problems can corrupt the software image.

- Connectivity (to retrieve images from remote servers)

- Configure the IP address for the 10/100BASE-T Ethernet port connection (interface mgmt0).
- Ensure the switch has a route to the remote server. The switch and the remote server must be in the same subnetwork if you do not have a router to route traffic between subnets.

Send documentation comments to mdsfeedback-doc@cisco.com.

- Images
 - Ensure that the specified system and kickstart images are compatible with each other.
 - If the kickstart image is not specified, the switch uses the current running kickstart image.
 - If you specify a different system image, ensure that it is compatible with the running kickstart image.
 - Retrieve images in one of two ways:
 - Locally—images are locally available on the switch.
 - Remotely—images are in a remote location and the user specifies the destination using the remote server parameters and the file name to be used locally.

- Terminology

[Table 6-1](#) summarizes terms used in this chapter with specific reference to the install and upgrade process.

Table 6-1 Terms Specific to This Chapter

| Term | | Definition |
|--------------|-----------|---|
| bootable | | The modules ability to boot or not boot based on image compatibility. |
| impact | | The type of software upgrade mechanism—disruptive or nondisruptive. |
| install-type | reset | Resets the module. |
| | sw-reset | Resets the module immediately after switchover. |
| | rolling | Upgrades each module in sequence. |
| | copy-only | Updates the software for BIOS, loader, or bootrom. |

- Commands
 - Verify connectivity to the remote server using the **ping** command.
 - Ensure that the required space is available for the image files to be copied using the **dir** command.
 - We recommend the one-step **install all** command to upgrade your software. This command upgrades all modules in any Cisco MDS 9000 Family switch (see the [“Benefits of Using the install all Command”](#) section on page 6-6).
 - Only one **install all** command can be running on a switch at any time.
 - No other command can be issued while running the **install all** command.
 - The **install all** command cannot be performed on the standby supervisor module—it can only be issued on the active supervisor module.
 - If the switching module(s) are not compatible with the new supervisor module image, some traffic disruption may be noticed in the related modules, depending on your configuration. These modules are identified in the summary when you issue the **install all** command. You can choose to proceed with the upgrade or end at this point.



Note

When you issue the **install all** command, the switch displays a summary of changes that are made to your configuration and waits for your authorization to continue executing the command process.

Send documentation comments to mdsfeedback-doc@cisco.com.

**Note**

To preserve the FC IDs in your configuration, verify that the persistent FC ID feature is enable before rebooting. As of SAN-OS Release 2.0(1b), this feature is enable by default. In earlier releases, the default is disabled. See the [“Persistent FC IDs” section on page 31-9](#).

Software Upgrade Methods

You can upgrade software without any disruptions using the Cisco MDS SAN-OS software designed for mission-critical high availability environments. To realize the benefits of nondisruptive upgrades on the Cisco MDS 9500 Directors, we highly recommend that you install dual supervisor modules.

You can upgrade any switch in the Cisco MDS 9000 Family using one of the following methods:

- Automated, one-step upgrade using the **install all** command. This upgrade is nondisruptive for directors in the Cisco MDS 9500 Series (see the [“Automated Upgrades” section on page 6-5](#)).
- Quick, one-step upgrade using the **reload** command. This upgrade is disruptive (see the [“Quick Upgrade” section on page 6-23](#)).

**Tip**

The **install all** command compares and presents the results of the compatibility before proceeding with the installation. You can exit if you do not want to proceed with these changes.

In some cases, regardless of which process you use, the software upgrades may be disruptive. These exception scenarios can occur under the following conditions:

- A single supervisor system with kickstart or system image changes.
- A dual supervisor system with incompatible system software images.

Determining Software Compatibility

If the running image and the image you want to install are incompatible, the software reports the incompatibility. In some cases, you may decide to proceed with this installation. If the active and the standby supervisor modules run different versions of the image, both images may be HA compatible in some cases and incompatible in others.

Compatibility is established based on the image and configuration:

- Image incompatibility—The running image and the image to be installed are not compatible.
- Configuration incompatibility—There is a possible compatibility if certain features in the running image are turned off as they are not supported in the image to be installed. The image to be installed is considered incompatible with the running image if one of the following statements is true:
 - An incompatible feature is enabled in the image to be installed and it is not available in the running image and may cause the switch to move into an inconsistent state. In this case, the incompatibility is *strict*.
 - An incompatible feature is enabled in the image to be installed and it is not available in the running image and does not cause the switch to move into an inconsistent state. In this case, the incompatibility is *loose*.

Send documentation comments to mdsfeedback-doc@cisco.com.

To view the results of a dynamic compatibility check, issue the **show incompatibility system bootflash:filename** command (see [Example 6-1](#)). Use this command to obtain further information when the **install all** command returns the following message:

Warning: The startup config contains commands not supported by the standby supervisor; as a result, some resources might become unavailable after a switchover.

Do you wish to continue? (y/ n) [y]: **n**

Example 6-1 Displays HA Compatibility Status

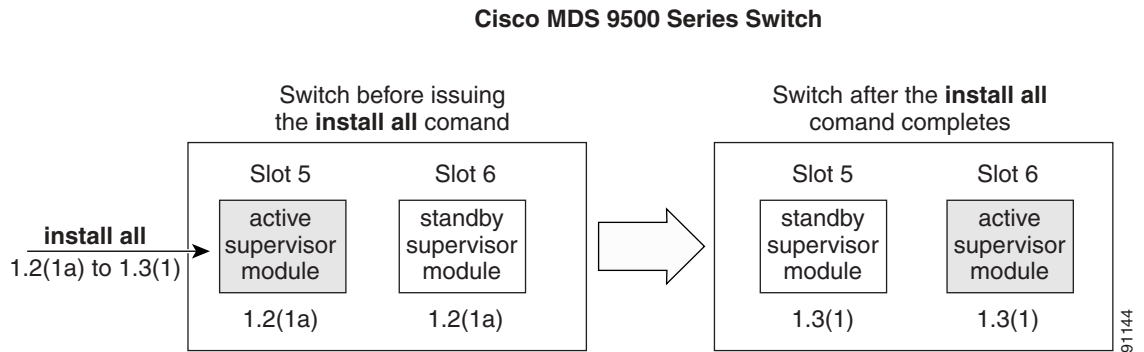
```
switch# show incompatibility system bootflash:running-image
The following configurations on active are incompatible with the system image
1) Feature Index : 67 , Capability : CAP_FEATURE_SPAN_FC_TUNNEL_CFG
Description : SPAN - Remote SPAN feature using fc-tunnels
Capability requirement : STRICT

2) Feature Index : 119 , Capability : CAP_FEATURE_FC_TUNNEL_CFG
Description : fc-tunnel is enabled
Capability requirement : STRICT
```

Automated Upgrades

The **install all** command upgrades all modules in any Cisco MDS 9000 Family switch. [Figure 6-1](#) provides an overview of the switch status before and after issuing the **install all** command.

Figure 6-1 The Effect of the install all Command



The **install all** command automatically verifies if the standby supervisor module is functioning (if present). If it is not functioning, it reloads that module and uses the **reload module slot force-dnld** command to force it to function.



Note

The **install all** command is only effective on switches running Cisco MDS SAN-OS Release 1.0(3) and later.

Send documentation comments to mdsfeedback-doc@cisco.com.

Benefits of Using the `install all` Command

The **`install all`** command provides the following benefits:

- You can upgrade the entire switch using just one command.
- You can receive descriptive information on the intended changes to your system before you continue with the command.
- You have the option to cancel the command. Once the effects of the command are presented, you can continue or cancel when you see this question (the default is **no**):

```
Do you want to continue (y/n) [n] :y
```
- You can upgrade the entire switch using the least disruptive procedure.
- You can see the progress of this command on the console, Telnet, and SSH screens:
 - After a switchover process, you can see the progress from both the supervisor modules.
 - Before a switchover process, you can only see the progress from the active supervisor module.
- The command automatically checks the image integrity. This includes the running kickstart and system images.
- The command performs a platform validity check to verify that a wrong image is not used—for example, to check if an MDS 9500 Series image is used inadvertently to upgrade an MDS 9200 Series switch.
- The **Ctrl-c** escape sequence gracefully ends the command. The command sequence completes the update step in progress and returns to the switch prompt. (Other upgrade steps cannot be ended using **Ctrl-c**.)
- After issuing the command, if any step in the sequence fails, the command completes the step in progress and ends.

For example, if a switching module fails to be updated for any reason (for example, due to an unstable fabric state), then the command sequence disruptively updates that module and ends. In such cases, you can verify the problem on the affected switching module and upgrade the other switching modules.

Recognizing Failure Cases

The following situations cause the **`install all`** command to end:

- If the standby supervisor module bootflash: filesystem does not have sufficient space to accept the updated image.
- If the specified system and kickstart images are not compatible.
- If the **`install all`** command is issued on the standby supervisor module.
- If the fabric or switch is configured while the upgrade is in progress.
- If a module is removed while the upgrade is in progress.
- If the switch has any power disruption while the upgrade is in progress.

Send documentation comments to mdsfeedback-doc@cisco.com.

- If the entire path for the remote location is not specified accurately.
- If images are incompatible after an upgrade. For example, a switching module image may be incompatible with the system image, or a kickstart image may be incompatible with a system image. This is also identified by the **show install all impact** command in the compatibility check section of the output (under the Bootable column).



Caution

If the **install all** command is ended, be sure to verify the state of the switch at every stage and reissue the command after 10 seconds. If you reissue the **install all** command within the 10-second span, the command is rejected with an error message indicating that an installation is currently in progress.



Tip

Most configurations are disallowed while the **install all** command is in progress. However, configurations coming through the CFS applications are allowed and may affect the upgrade procedure.

Using the install all Command



Note

Ensure that there is enough space available on the active and standby supervisor bootflash to store the images being installed, even if the images are supplied in slot0. The system will automatically sync the images to the standby supervisor.

To perform an automated software upgrade on any switch, follow these steps:

- Step 1** Log into the switch through the console, Telnet, or SSH port of the active supervisor.
- Step 2** Create a backup of your existing configuration file, if required (see the [“Working with Configuration Files”](#) section on page 4-24).
- Step 3** Perform the upgrade by issuing the **install all** command.

```
switch# install all system bootflash:isan-1.3.1 kickstart bootflash:boot-1.3.1
```

```
Verifying image bootflash:/boot-1.3.1
[#####] 100% -- SUCCESS
```

```
Verifying image bootflash:/isan-1.3.1
[#####] 100% -- SUCCESS
```

```
Extracting "slc" version from image bootflash:/isan-1.3.1.
[#####] 100% -- SUCCESS
```

```
Extracting "ips" version from image bootflash:/isan-1.3.1.
[#####] 100% -- SUCCESS
```

```
Extracting "system" version from image bootflash:/isan-1.3.1.
[#####] 100% -- SUCCESS
```

```
Extracting "kickstart" version from image bootflash:/boot-1.3.1.
[#####] 100% -- SUCCESS
```

```
Extracting "loader" version from image bootflash:/boot-1.3.1.
[#####] 100% -- SUCCESS
```

■ Automated Upgrades

| Module | bootable | Impact | Install-type | Reason |
|--------|----------|--------|--------------|--------|
|--------|----------|--------|--------------|--------|

| Module | bootable | Impact | Install-type | Reason |
|--------|----------|--------|--------------|--------|
|--------|----------|--------|--------------|--------|

| Module | Image | Running-Version | New-Version | Upq-Required |
|--------|-------|-----------------|-------------|--------------|
|--------|-------|-----------------|-------------|--------------|

| Module | Image | Running-Version | New-Version | Upq-Required |
|--------|-------|-----------------|-------------|--------------|
|--------|-------|-----------------|-------------|--------------|

Install is in progress, please wait.

Syncing image bootflash:/boot-1.3.1 to standby.

```
[#####] 100% -- SUCCESS
```

Syncing image bootflash:/isan-1.3.1 to standby.

```
[#####] 100% -- SUCCESS
```

```
Jan 18 23:40:03 Hacienda %VSHD-5-VSHD SYSLOG CONFIG I: Configuring console from
```

Performing configuration copy.

```
[#####] 100% -- SUCCESS
```

Module 6: Waiting for module online.

1

```
Auto booting bootflash:/boot-1.3.1 bootflash:/isan-1.3.1...
```

```
Booting kickstart image: bootflash:/boot-1.3.1...
```

```
.....Image verification OK
```

Starting kernel...

```
INIT: version 2.78 booting
```

```
Checking all filesystems..r.r.. done.
```

Loading system software

```
Uncompressing system image: bootflash:/isan-1.3.1
```

[illegible]

```
INIT: Entering runlevel: 3
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Step 4 Exit the switch console and open a new terminal session to view the upgraded supervisor module using the **show module** command.

If the configuration meets all guidelines when the **install all** command is issued, all modules (supervisor and switching) are upgraded. This is true for any switch in the Cisco MDS 9000 Family.

**Caution**

If a nondisruptive upgrade operation fails for any reason other than those listed in the “[Recognizing Failure Cases](#)” section on page 6-6, contact your reseller or Cisco representative for further assistance.

If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco Systems, contact Cisco Technical Support at this URL: <http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

Upgrading Services Modules

Any Fibre Channel switching module supports nondisruptive upgrades. The 14/2-port Multiprotocol Services (MPS-14/2)) module supports nondisruptive upgrades for the Fibre Channel ports. Any software upgrade for the two Gigabit Ethernet ports in this module is disruptive. Refer to the [Chapter 28](#), “[Configuring IP Storage](#)” for further information on MPS-14/2 modules.

**Caution**

Any software upgrade for the Caching Services Module (CSM) and the IP Storage (IPS) services modules is disruptive.

CSMs and IPS modules use a rolling upgrade install mechanism to guarantee a stable state for each module in the switch:

- Each IPS module in a switch requires a 5-minute delay before the next IPS module is upgraded. Refer to the [Chapter 28](#), “[Configuring IP Storage](#)” for further information on IPS modules.
- Each CSM module requires a 30-minute delay before the next CSM module is upgraded. Refer to the *Cisco MDS 9000 Family SAN Volume Controller Configuration Guide* for more information on CSMs.

Sample install all Commands

[Example 6-2](#) displays the result of the **install all** command issued from a console terminal that is connected to the active supervisor. Once a switchover happens, you can see the rest of the output from the console terminal of the standby supervisor module. [Example 6-3](#) displays the file output continuation of the **install all** command on the console of the standby supervisor module. [Example 6-4](#) displays the result of the **install all** command issued from a console terminal for a system that contains an SSI image.

Similarly, you can view the results of the **install all** command issued from the SSH or Telnet terminal that is connected to the active supervisor. Once a switchover happens, you need to log back into the switch and issue the **show install all status** command (see the “[Upgrade Status Verification](#)” section on page 6-17).

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 6-2 Successful install all Command Issued from the Active Console

```
Hacienda# install all system bootflash:isan-1.3.1 kickstart bootflash:boot-1.3.1
```

```
Verifying image bootflash:/boot-1.3.1
[#####] 100% -- SUCCESS
```

```
Verifying image bootflash:/isan-1.3.1
[#####] 100% -- SUCCESS
```

```
Extracting "slc" version from image bootflash:/isan-1.3.1.
[#####] 100% -- SUCCESS
```

```
Extracting "ips" version from image bootflash:/isan-1.3.1.
[#####] 100% -- SUCCESS
```

```
Extracting "system" version from image bootflash:/isan-1.3.1.
[#####] 100% -- SUCCESS
```

```
Extracting "kickstart" version from image bootflash:/boot-1.3.1.
[#####] 100% -- SUCCESS
```

```
Extracting "loader" version from image bootflash:/boot-1.3.1.
[#####] 100% -- SUCCESS
```

Compatibility check is done:

| Module | bootable | Impact | Install-type | Reason |
|--------|----------|----------------|--------------|----------------------------------|
| 1 | yes | non-disruptive | rolling | |
| 2 | yes | disruptive | rolling | Hitless upgrade is not supported |
| 3 | yes | disruptive | rolling | Hitless upgrade is not supported |
| 4 | yes | non-disruptive | rolling | |
| 5 | yes | non-disruptive | reset | |
| 6 | yes | non-disruptive | reset | |

Images will be upgraded according to following table:

| Module | Image | Running-Version | New-Version | Upg-Required |
|--------|-----------|------------------|------------------|--------------|
| 1 | slc | 1.3(2a) | 1.3(1) | yes |
| 1 | bios | v1.1.0(10/24/03) | v1.1.0(10/24/03) | no |
| 2 | ips | 1.3(2a) | 1.3(1) | yes |
| 2 | bios | v1.1.0(10/24/03) | v1.1.0(10/24/03) | no |
| 3 | ips | 1.3(2a) | 1.3(1) | yes |
| 3 | bios | v1.1.0(10/24/03) | v1.1.0(10/24/03) | no |
| 4 | slc | 1.3(2a) | 1.3(1) | yes |
| 4 | bios | v1.1.0(10/24/03) | v1.1.0(10/24/03) | no |
| 5 | system | 1.3(2a) | 1.3(1) | yes |
| 5 | kickstart | 1.3(2a) | 1.3(1) | yes |
| 5 | bios | v1.1.0(10/24/03) | v1.1.0(10/24/03) | no |
| 5 | loader | 1.2(2) | 1.2(2) | no |
| 6 | system | 1.3(2a) | 1.3(1) | yes |
| 6 | kickstart | 1.3(2a) | 1.3(1) | yes |
| 6 | bios | v1.1.0(10/24/03) | v1.1.0(10/24/03) | no |
| 6 | loader | 1.2(2) | 1.2(2) | no |

Do you want to continue with the installation (y/n)? [n] **y**

Install is in progress, please wait.

```
Syncing image bootflash:/boot-1.3.1 to standby.  
[#####] 100% -- SUCCESS  
  
Syncing image bootflash:/isan-1.3.1 to standby.  
[#####] 100% -- SUCCESS  
Jan 18 23:40:03 Hacienda %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configuring console from  
  
Performing configuration copy.  
[#####] 100% -- SUCCESS  
  
Module 6: Waiting for module online.  
|  
Auto booting bootflash:/boot-1.3.1 bootflash:/isan-1.3.1...  
Booting kickstart image: bootflash:/boot-1.3.1....  
.....Image verification OK  
  
Starting kernel...  
INIT: version 2.78 booting  
Checking all filesystems..r.r.. done.  
Loading system software  
Uncompressing system image: bootflash:/isan-1.3.1  
CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC  
CCCCCCCCCCCCCCCCCCCCCCCCCCC  
INIT: Entering runlevel: 3
```

Example 6-3 Successful install all Command Output Continued from the Standby Console

Cisco MDS 9000 Family Configuration Guide

Send documentation comments to mdsfeedback-doc@cisco.com.

```
Jan 18 23:44:54 Hacienda %BOOTVAR-5-NEIGHBOR_UPDATE_AUTOCOPY: auto-copy supported by
neighbor, starting...
Module 1: Non-disruptive upgrading.
[#          ] 0%Jan 18 23:44:56 Hacienda %MODULE-5-STANDBY_SUP_OK: Supervisor 5
is standby
Jan 18 23:44:55 Hacienda %IMAGE_DNLD-SLOT1-2-IMG_DNLD_STARTED: Module image download
process. Please wait until completion...
Jan 18 23:45:12 Hacienda %IMAGE_DNLD-SLOT1-2-IMG_DNLD_COMPLETE: Module image download
process. Download successful.
Jan 18 23:45:48 Hacienda %MODULE-5-MOD_OK: Module 1 is online
[#####] 100% -- SUCCESS
Module 4: Non-disruptive upgrading.
[#          ] 0%Jan 18 23:46:12 Hacienda %IMAGE_DNLD-SLOT4-2-IMG_DNLD_STARTED:
Module image download process. Please wait until completion...
Jan 18 23:46:26 Hacienda %IMAGE_DNLD-SLOT4-2-IMG_DNLD_COMPLETE: Module image download
process. Download successful.
Jan 18 23:47:02 Hacienda %MODULE-5-MOD_OK: Module 4 is online
[#####] 100% -- SUCCESS
Module 2: Disruptive upgrading.
...
-- SUCCESS
Module 3: Disruptive upgrading.
...
-- SUCCESS
Install has been successful.
MDS Switch
Hacienda login:
```

Example 6-4 displays the result of the **install all** command issued from a console terminal for a system that contains an SSI image. The **install all** command syncs the SSI image to the standby supervisor.



Note

You can use the **install all** command for the SSM only if the SSM is already up and running. For first time SSM installations, see the “[Configuring the SSI Image Boot Variable for Fibre Channel Switching and Intelligent Storage Services](#)” section on page 7-21.

Example 6-4 Successful install all Command Including an SSI Image

```
Cisco-MDS# install all system bootflash:isan-2-1-1a kickstart
bootflash:boot-2-1-1a ssi bootflash:ssi-2.1.1a

Verifying image bootflash:/ssi-2.1.1a
[#####] 100% -- SUCCESS

Verifying image bootflash:/boot-2-1-1a
[#####] 100% -- SUCCESS

Verifying image bootflash:/isan-2-1-1a
[#####] 100% -- SUCCESS

Extracting "slc" version from image bootflash:/isan-2-1-1a.
[#####] 100% -- SUCCESS

Extracting "ips4" version from image bootflash:/isan-2-1-1a.
[#####] 100% -- SUCCESS

Extracting "system" version from image bootflash:/isan-2-1-1a.
[#####] 100% -- SUCCESS

Extracting "kickstart" version from image bootflash:/boot-2-1-1a.
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
[#####] 100% -- SUCCESS
```

```
Extracting "loader" version from image bootflash:/boot-2-1-1a.
```

```
[#####] 100% -- SUCCESS
```

Compatibility check is done:

| Module | bootable | Impact | Install-type | Reason |
|--------|----------|----------------|--------------|----------------------------------|
| 2 | yes | non-disruptive | rolling | |
| 3 | yes | disruptive | rolling | Hitless upgrade is not supported |
| 4 | yes | disruptive | rolling | Hitless upgrade is not supported |
| 5 | yes | non-disruptive | reset | |

Images will be upgraded according to following table:

| Module | Image | Running-Version | New-Version | Upg-Required |
|--------|-----------|------------------|------------------|--------------|
| 2 | slc | 2.0(3) | 2.1(1a) | yes |
| 2 | bios | v1.1.0(10/24/03) | v1.1.0(10/24/03) | no |
| 3 | slc | 2.0(3) | 2.1(1a) | yes |
| 3 | ssi | 2.0(3) | 2.1(1a) | yes |
| 3 | bios | v1.0.8(08/07/03) | v1.1.0(10/24/03) | yes |
| 4 | ips4 | 2.0(3) | 2.1(1a) | yes |
| 4 | bios | v1.1.0(10/24/03) | v1.1.0(10/24/03) | no |
| 5 | system | 2.0(3) | 2.1(1a) | yes |
| 5 | kickstart | 2.0(3) | 2.1(1a) | yes |
| 5 | bios | v1.1.0(10/24/03) | v1.1.0(10/24/03) | no |
| 5 | loader | 1.2(2) | 1.2(2) | no |

Do you want to continue with the installation (y/n)? [n] y

Install is in progress, please wait.

Module 6:Force downloading.

```
-- SUCCESS
```

Syncing image bootflash:/ssi-2.1.1a to standby.

```
[#####] 100% -- SUCCESS
```

Syncing image bootflash:/boot-2-1-1a to standby.

```
[#####] 100% -- SUCCESS
```

Syncing image bootflash:/isan-2-1-1a to standby.

```
[#####] 100% -- SUCCESS
```

Setting boot variables.

```
[#####] 100% -- SUCCESS
```

Performing configuration copy.

```
[#####] 100% -- SUCCESS
```

Module 3:Upgrading Bios/loader/bootrom.

```
[#####] 100% -- SUCCESS
```

Module 6:Waiting for module online.

```
-- SUCCESS
```

"Switching over onto standby".

```
-----
```

Send documentation comments to mdsfeedback-doc@cisco.com.



Note

If you perform the **install all** command to downgrade to a Cisco MDS SAN-OS release that does not support the SSM module, you must power down the SSM module when prompted by the CLI console. The boot variables for the SSM module are lost.

Example 6-5 displays the result of the **install all** command if the system and kickstart files are automatically downloaded using a remote (TFTP, FTP, SCP, or SFTP) download option. It shows an accurate and complete example.



Caution

Specify the complete path of the remote location. The system will not allow you to proceed if the entire path is not accurately specified. Here are examples of incomplete **install all** commands.

```
switch# install all system bootflash:system-image kickstart tftp:
Please provide a complete URI
switch# install all system scp:
Please provide a complete URI
```

Example 6-5 A Sample of the install all Command Issued Using a Remote Download

```
switch# install all system
scp://user@171.69.16.26/tftpboot/HKrel/qa/final/m9500-sflek9-mz.1.3.2a.bin kickstart
scp://user@171.69.16.26/tftpboot/HKrel/qa/final/m9500-sflek9-kickstart-mz.1.3.2a.bin
For scp://user@171.69.16.26, please enter password:
For scp://user@171.69.16.26, please enter password:

Copying image from
scp://user@171.69.16.26/tftpboot/HKrel/qa/final/m9500-sflek9-kickstart-mz.1.3.2a.bin to
bootflash:///m9500-sflek9-kickstart-mz.1.3.2a.bin.
[#####] 100% -- SUCCESS

Copying image from
scp://user@171.69.16.26/tftpboot/HKrel/qa/final/m9500-sflek9-mz.1.3.2a.bin to
bootflash:///m9500-sflek9-mz.1.3.2a.bin.
[#####] 100% -- SUCCESS

Verifying image bootflash:///m9500-sflek9-kickstart-mz.1.3.2a.bin
[#####] 100% -- SUCCESS

Verifying image bootflash:///m9500-sflek9-mz.1.3.2a.bin
[#####] 100% -- SUCCESS

Extracting "slc" version from image bootflash:///m9500-sflek9-mz.1.3.2a.bin.
[#####] 100% -- SUCCESS

Extracting "ips" version from image bootflash:///m9500-sflek9-mz.1.3.2a.bin.
[#####] 100% -- SUCCESS

Extracting "system" version from image bootflash:///m9500-sflek9-mz.1.3.2a.bin.
[#####] 100% -- SUCCESS

Extracting "kickstart" version from image
bootflash:///m9500-sflek9-kickstart-mz.1.3.2a.bin.
[#####] 100% -- SUCCESS

Extracting "loader" version from image bootflash:///m9500-sflek9-kickstart-mz.1.3.2a.bin.
[#####] 100% -- SUCCESS
```


Send documentation comments to mdsfeedback-doc@cisco.com.

Compatibility check is done:

| Module | bootable | Impact | Install-type | Reason |
|--------|----------|----------------|--------------|----------------------------------|
| ----- | ----- | ----- | ----- | ----- |
| 1 | yes | non-disruptive | rolling | |
| 2 | yes | disruptive | rolling | Hitless upgrade is not supported |
| 3 | yes | non-disruptive | rolling | |
| 4 | yes | non-disruptive | rolling | |
| 5 | yes | non-disruptive | reset | |
| 6 | yes | non-disruptive | reset | |
| 7 | yes | non-disruptive | rolling | |
| 8 | yes | non-disruptive | rolling | |
| 9 | yes | disruptive | rolling | Hitless upgrade is not supported |

Images will be upgraded according to following table:

| Module | Image | Running-Version | New-Version | Upg-Required |
|--------|-----------|------------------|------------------|--------------|
| ----- | ----- | ----- | ----- | ----- |
| 1 | slc | 1.3(1) | 1.3(2a) | yes |
| 1 | bios | v1.1.0(10/24/03) | v1.0.8(08/07/03) | no |
| 2 | ips | 1.3(1) | 1.3(2a) | yes |
| 2 | bios | v1.1.0(10/24/03) | v1.0.8(08/07/03) | no |
| 3 | slc | 1.3(1) | 1.3(2a) | yes |
| 3 | bios | v1.1.0(10/24/03) | v1.0.8(08/07/03) | no |
| 4 | slc | 1.3(1) | 1.3(2a) | yes |
| 4 | bios | v1.1.0(10/24/03) | v1.0.8(08/07/03) | no |
| 5 | system | 1.3(1) | 1.3(2a) | yes |
| 5 | kickstart | 1.3(1) | 1.3(2a) | yes |
| 5 | bios | v1.1.0(10/24/03) | v1.0.8(08/07/03) | no |
| 5 | loader | 1.2(2) | 1.2(2) | no |
| 6 | system | 1.3(1) | 1.3(2a) | yes |
| 6 | kickstart | 1.3(1) | 1.3(2a) | yes |
| 6 | bios | v1.1.0(10/24/03) | v1.0.8(08/07/03) | no |
| 6 | loader | 1.2(2) | 1.2(2) | no |
| 7 | slc | 1.3(1) | 1.3(2a) | yes |
| 7 | bios | v1.1.0(10/24/03) | v1.0.8(08/07/03) | no |
| 8 | slc | 1.3(1) | 1.3(2a) | yes |
| 8 | bios | v1.1.0(10/24/03) | v1.0.8(08/07/03) | no |
| 9 | ips | 1.3(1) | 1.3(2a) | yes |
| 9 | bios | v1.1.0(10/24/03) | v1.0.8(08/07/03) | no |

Do you want to continue with the installation (y/n)? [n]

Example 6-6 displays the **install all** command output of a failed operation due to a lack of disk space.

Example 6-6 Failed Operation Due to a Full bootflash: File System

```
switch# install all system bootflash:isan-1.3.2a kickstart bootflash:boot-1.3.2a
```

```
Verifying image bootflash:/boot-1.3.2a
[#####] 100% -- SUCCESS
```

```
Verifying image bootflash:/isan-1.3.2a
[#####] 100% -- SUCCESS
```

```
Extracting "slc" version from image bootflash:/isan-1.3.2a.
[#####] 100% -- SUCCESS
```

```
Extracting "ips" version from image bootflash:/isan-1.3.2a.
[#####] 100% -- SUCCESS
```

```
Extracting "system" version from image bootflash:/isan-1.3.2a.
[#####] 100% -- SUCCESS
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
Extracting "kickstart" version from image bootflash:/boot-1.3.2a.
[#####] 100% -- SUCCESS
```

```
Extracting "loader" version from image bootflash:/boot-1.3.2a.
[#####] 100% -- SUCCESS
```

Compatibility check is done:

| Module | bootable | Impact | Install-type | Reason |
|--------|----------|----------------|--------------|----------------------------------|
| 1 | yes | non-disruptive | rolling | |
| 2 | yes | disruptive | rolling | Hitless upgrade is not supported |
| 3 | yes | non-disruptive | rolling | |
| 4 | yes | non-disruptive | rolling | |
| 5 | yes | non-disruptive | reset | |
| 6 | yes | non-disruptive | reset | |
| 7 | yes | non-disruptive | rolling | |
| 8 | yes | non-disruptive | rolling | |
| 9 | yes | disruptive | rolling | Hitless upgrade is not supported |

Images will be upgraded according to following table:

| Module | Image | Running-Version | New-Version | Upg-Required |
|--------|-----------|------------------|------------------|--------------|
| 1 | slc | 1.3(1) | 1.3(2a) | yes |
| 1 | bios | v1.1.0(10/24/03) | v1.0.8(08/07/03) | no |
| 2 | ips | 1.3(1) | 1.3(2a) | yes |
| 2 | bios | v1.1.0(10/24/03) | v1.0.8(08/07/03) | no |
| 3 | slc | 1.3(1) | 1.3(2a) | yes |
| 3 | bios | v1.1.0(10/24/03) | v1.0.8(08/07/03) | no |
| 4 | slc | 1.3(1) | 1.3(2a) | yes |
| 4 | bios | v1.1.0(10/24/03) | v1.0.8(08/07/03) | no |
| 5 | system | 1.3(1) | 1.3(2a) | yes |
| 5 | kickstart | 1.3(1) | 1.3(2a) | yes |
| 5 | bios | v1.1.0(10/24/03) | v1.0.8(08/07/03) | no |
| 5 | loader | 1.2(2) | 1.2(2) | no |
| 6 | system | 1.3(1) | 1.3(2a) | yes |
| 6 | kickstart | 1.3(1) | 1.3(2a) | yes |
| 6 | bios | v1.1.0(10/24/03) | v1.0.8(08/07/03) | no |
| 6 | loader | 1.2(2) | 1.2(2) | no |
| 7 | slc | 1.3(1) | 1.3(2a) | yes |
| 7 | bios | v1.1.0(10/24/03) | v1.0.8(08/07/03) | no |
| 8 | slc | 1.3(1) | 1.3(2a) | yes |
| 8 | bios | v1.1.0(10/24/03) | v1.0.8(08/07/03) | no |
| 9 | ips | 1.3(1) | 1.3(2a) | yes |
| 9 | bios | v1.1.0(10/24/03) | v1.0.8(08/07/03) | no |

Do you want to continue with the installation (y/n)? [n] y

Install is in progress, please wait.

```
Syncing image bootflash:/boot-1.3.2a to standby.
[#####] 100% -- SUCCESS
```

```
Syncing image bootflash:/isan-1.3.2a to standby.
[#          ] 0% -- FAIL. Return code 0x401E0008 (request was aborted, standby
disk may be full).
```

```
Install has failed. Return code 0x40930013 (Syncing images to standby failed). <----
-----insufficient space message
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Please identify the cause of the failure, and try 'install all' again.
Dec 15 19:36:42 switch %SYSMGR-3-SERVICE_TERMINATED: Service "installer" (PID 5470) has finished with error code SYSMGR_EXITCODE_FAILURE_NOCALLHOME (20).

[Example 6-7](#) displays the **install all** command output of a failed operation due to an invalid image.

Example 6-7 Failed Operation Due to an Invalid Image

```
install all system bootflash:junk kickstart bootflash:junk

Verifying image bootflash:/junk
[#          ] 0% -- FAIL. Return code 0x4045001E (mismatch between actual image
type and boot variable).
Compatibility check failed. Return code 0x40930011 (Image verification failed).
Hacienda# Jan 19 00:20:35 Hacienda %SYSMGR-3-SERVICE_TERMINATED: Service "installer" (PID
5664) has finished with error code SYSMGR_EXITCODE_FAILURE_NOCALLHOME (20).
```

Upgrade Status Verification

Use the **show install all status** command to view the on-going **install all** command or the log of the last installed **install all** command from a console, SSH, or Telnet session.

This command presents the **install all** output on both the active and standby supervisor module even if you are not connected to the console terminal. It only displays the status of an **install all** command that is issued from the CLI (not the GUI). See [Example 6-8](#).

Example 6-8 Displays the install all Command Output

```
switch# show install all status
There is an on-going installation... <----- in progress installation
Enter Ctrl-C to go back to the prompt.
Verifying image bootflash:/b-1.3.0.104
-- SUCCESS

Verifying image bootflash:/i-1.3.0.104
-- SUCCESS

Extracting "system" version from image bootflash:/i-1.3.0.104.
-- SUCCESS

Extracting "kickstart" version from image bootflash:/b-1.3.0.104.
-- SUCCESS

Extracting "loader" version from image bootflash:/b-1.3.0.104.
-- SUCCESS

switch# show install all status
This is the log of last installation.          <<<<<< log of last install

Verifying image bootflash:/b-1.3.0.104
-- SUCCESS

Verifying image bootflash:/i-1.3.0.104
-- SUCCESS

Extracting "system" version from image bootflash:/i-1.3.0.104.
-- SUCCESS
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
Extracting "kickstart" version from image bootflash:/b-1.3.0.104.
-- SUCCESS

Extracting "loader" version from image bootflash:/b-1.3.0.104.
-- SUCCESS
```

Manual Upgrade on a Dual Supervisor Switch



Caution

If you are a new user, use the **install all** command to perform a software upgrade. This section is for administrators or individuals who are completely familiar with specific switch functions.

You can manually upgrade the BIOS and the loader in any Cisco MDS switch using the procedures provided in this section. This upgrade process requires you to implement some or all procedures depending on your switch or network configuration.

To perform a manual upgrade for a dual supervisor switch, follow these steps.

-
- Step 1** [“Preparing for a Manual Installation” section on page 6-18](#)
 - Step 2** [“Upgrading a Loader” section on page 6-19](#)
 - Step 3** [“Upgrading the BIOS” section on page 6-21](#)
-

Preparing for a Manual Installation

To prepare any Cisco MDS 9000 Family switch for a manual software installation, follow these steps:

-
- Step 1** Log into the switch through the console port, an SSH session, or a Telnet session.
 - Step 2** Create a backup of your existing configuration file, if required (see the [“Saving the Configuration” section on page 4-27](#)).
 - Step 3** Copy the software image from a SCP location to one of two targets: bootflash: or slot0:.
- The switch remains operational while the image file is copied.
- Bootflash device (SCP defaults to the bootflash device)—Copy the software image file from the appropriate SCP file system to the bootflash: file system.

```
switch# copy scp://server_IP_address/destination_file_name
```

For example:

```
switch# copy scp://user@10.1.7.2/system-image bootflash:system-image
```



Note

The Cisco MDS 9216 Switch does not have an external CompactFlash (see the [“Working with Configuration Files” section on page 4-24](#)). If you are using a switch in this series, use the bootflash: file system to copy and verify files.

Send documentation comments to mdsfeedback-doc@cisco.com.

- CompactFlash device—Copy the software image file from the appropriate SCP file system to the CompactFlash device in slot0: file system.

```
switch# copy scp://server_IP_address/file_name_in_SCP slot0:system-image
```

You can also copy the image onto a new Flash disk from a PC and insert it in slot0: in the Cisco MDS 9500 Series switch. After you copy the image and insert it into the slot0: file system, the process is the same as the CompactFlash device after the **copy** command is issued.

Step 4 Verify that the file was copied in the required directory.

```
switch# dir bootflash:
40295206      Aug 05 15:23:51 1980  ilc1.bin
12456448      Jul 30 23:05:28 1980  kickstart-image1
12288         Jun 23 14:58:44 1980  lost+found/
27602159      Jul 30 23:05:16 1980  system-image1
12447232      Aug 05 15:08:30 1980  kickstart-image2
28364853      Aug 05 15:11:57 1980  system-image2
```

```
Usage for bootflash://sup-local
135404544 bytes used
49155072 bytes free
184559616 bytes total
```

Step 5 Ensure that the software images are not damaged or corrupted in the saved bootflash: file system. When copying a new image to your switch, confirm that the image was not corrupted during the copy process.

Use the **show version image** command to verify that the required image was copied successfully.

```
switch# show version image bootflash:kickstart-image
image name: m9500-sflek9-kickstart-mzg.1.0.3.bin
kickstart:  version 1.0(3)
loader:      version 1.0(3)
compiled:    2/12/2003 11:00:00
```



Note

A verification failed message is generated when you use a Cisco MDS 9500 Series image on a Cisco MDS 9200 Series switch or a Cisco MDS 9200 Series image on a Cisco MDS 9500 Series switch. Be sure to verify the right image.

Step 6 Compare the running system image and the new image by issuing the **show install all impact** command.

Upgrading a Loader

The **install module slot# of the supervisor module loader** command upgrades the (boot) loader.



Note

If the loader is upgraded, you need to reboot to make the new loader effective. You can schedule the reboot at a convenient time so traffic is not impacted.



Caution

Before issuing this command, be sure to read the release notes to verify compatibility issues between the loader and the kickstart or system images.

Send documentation comments to mdsfeedback-doc@cisco.com.

To upgrade the loader on either the active or standby supervisor module, follow these steps.

- Step 1** Use the **show version** command to verify the version on the active and standby supervisor modules.

```
switch# show version
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2003, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
Cisco Systems, Inc. and/or other third parties and are used and
distributed under license. Some parts of this software are covered
under the GNU Public License. A copy of the license is available
at http://www.gnu.org/licenses/gpl.html.

Software
  BIOS:          version 1.0.8
  loader:        version 1.1(2) <-----current running version
  kickstart:     version 2.0(1)
  system:        version 2.0(1)

  BIOS compile time:      08/07/03
  kickstart image file is: bootflash:///m9500-sflek9-kickstart-mzg.2.0.0.6.bin
  kickstart compile time: 10/25/2010 12:00:00
  system image file is:   bootflash:///m9500-sflek9-mzg.2.0.0.6.bin
  system compile time:    10/25/2020 12:00:00

Hardware
  RAM 1024584 kB

  bootflash: 1000944 blocks (block size 512b)
  slot0:      0 blocks (block size 512b)

172.22.92.181 uptime is 0 days 2 hours 18 minute(s) 1 second(s)

Last reset at 970069 usecs after Tue Sep 16 22:31:25 1980
Reason: Reset Requested by CLI command reload
System version: 2.0(0.6)
Service:
```

- Step 2** Issue the **install module** command for the required supervisor module (active or standby). This example displays the command being issued for the standby supervisor module in slot 6.

```
switch# install module 6 loader bootflash:kickstart-image
```



Note If you install a loader version that is the same as the currently installed version, the command will not execute. When both the current version and the installed version are the same, use the **init system** command to force a loader upgrade.

- Step 3** Use the **show version** command to verify the updated image on the supervisor module.

```
switch# show version
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2003, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
Cisco Systems, Inc. and/or other third parties and are used and
distributed under license. Some parts of this software are covered
under the GNU Public License. A copy of the license is available
at http://www.gnu.org/licenses/gpl.html.
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
Software
  BIOS:      version 1.3.1
  loader:    version 1.2(2) <-----New running version
  kickstart: version 1.3(1) ]
  system:    version 1.3(1)

  BIOS compile time:      08/07/03
  kickstart image file is: bootflash:///m9500-sflek9-kickstart-mzg.2.0.0.6.bin
  kickstart compile time: 10/25/2010 12:00:00
  system image file is:   bootflash:///m9500-sflek9-mzg.2.0.0.6.bin
  system compile time:    10/25/2020 12:00:00

Hardware
  RAM 1024584 kB

  bootflash: 1000944 blocks (block size 512b)
  slot0:      0 blocks (block size 512b)

  172.22.92.181 uptime is 0 days 2 hours 18 minute(s) 1 second(s)

  Last reset at 970069 usecs after Tue Sep 16 22:31:25 1980
  Reason: Reset Requested by CLI command reload
  System version: 2.0(0.6)
  Service:
```

Upgrading the BIOS



Tip

Refer to the release notes to verify if the BIOS has changed for the image version being used.

Program the supervisor or switching module BIOS only if a new BIOS image is provided by Cisco System. Only use the provided image to upgrade the BIOS. This command does not affect traffic and can be issued at any time on any switch in the Cisco MDS 9200 Series or Cisco MDS 9500 Series.



Note

If the BIOS is upgraded, reboot to make the new BIOS take effect. You can schedule the reboot at a convenient time so traffic is not impacted.



Caution

The console baud rate automatically reverts to the default rate (9600) after any BIOS upgrade.

Send documentation comments to mdsfeedback-doc@cisco.com.

To upgrade the BIOS for a module, follow these steps:

Step 1 Use the **show version** command to verify the current running BIOS version.

```
switch# show version
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2003 by Cisco Systems, Inc. All rights reserved.
The copyright for certain works contained herein are owned by
Cisco Systems, Inc. and/or other third parties and are used and
distributed under license.
Software
  BIOS:      version 1.0(6) <----- current running version
  loader:    version 1.0(3)
  kickstart: version 1.0(3)
  system:    version 1.0(3)

  BIOS compile time:      01/27/03
  kickstart image file is: bootflash:/kickstart-image
  kickstart compile time: 01/25/2003 12:00:00
  system image file is:   bootflash:/system-image
  system compile time:    01/25/2003 12:00:00

Hardware
  RAM 1027564 kB
```

Step 2 Verify that the BIOS version of the system image is different from the running image.

```
switch# show version image bootflash:system-image
  image name: m9500-sflek9-mz.1.0.3.bin
  bios:      version v1.0.6(01/27/03) <----- BIOS is same version 1.0.6
  system:    version 1.0(3)
  compiled:  2/28/2003 5:00:00

system service's list

package name      package version
acl               1.0(3)
ascii-cfg         1.0(3)
bios_daemon       1.0(3)
...
```



Note If the versions are different, issue the **install module** command as specified in Step 3. If they are the same, you do not need to update the BIOS image.

Send documentation comments to mdsfeedback-doc@cisco.com.

- Step 3** Run the **install module slot# bios** command to install each module (if required). In this example, the supervisor module in slot 6 was updated.

```
switch# install module 6 bios system bootflash:system-image
Started bios programming .... please wait
[#####] 100%
BIOS upgrade succeeded for module 1
```



Caution Do not reboot the switch if any errors were indicated in response to this command.

- Step 4** Issue the **show version** command to verify that the module was updated with a the new BIOS version.

```
switch# show version module 6
ModNo  Image Type  SW Version  SW Interim Version  BIOS Version
6      Stby Sup    1.3 (2)     1.3 (1.1)          1.1.0 [last 1.0.6]
```

Quick Upgrade

To perform a quick upgrade on a Cisco MDS 9000 Family switch, follow these steps:

- Step 1** Copy the kickstart and system image files to the required location (see the “Copying Files” section on [page 4-28](#)).
- Step 2** Set the boot variables.
- Step 3** Issue the **reload** command. The **reload** command reboots the system. This upgrade is disruptive.



Tip Use the **install all** command to gracefully reload the switch and handle configuration conversions.

Maintaining Supervisor Modules

This section includes general information about replacing and using supervisor modules effectively.

Standby Supervisor Boot Variable Version

If the standby supervisor module boot variable images are not the *same* version as those running on the active supervisor module, the software forces the standby supervisor module to run the same version as the active supervisor module.

If you specifically set the boot variables of the standby supervisor module to a different version and reboot the standby supervisor module, the standby supervisor module will only load the specified boot variable if the same version is also running on the active supervisor module. At this point, the standby supervisor module is *not* running the images set in the boot variables.

Send documentation comments to mdsfeedback-doc@cisco.com.

Standby Supervisor Bootflash Memory

When updating software images on the standby supervisor module, verify that there is enough space available for the image using the **dir bootflash://sup-standby/** command. It is a good practice to remove older versions of Cisco MDS SAN-OS images and kickstart images. For information about displaying file systems and removing files, see the [“Using the File System” section on page 2-22](#).

Standby Supervisor Boot Alert

If a standby supervisor module fails to boot, the active supervisor module detects that condition and generates a Call Home event and a system message and reboots the standby supervisor module approximately 3 to 6 minutes after the standby supervisor module moves to the `loader>` prompt.

The following system message is issued:

```
%DAEMON-2-SYSTEM_MSG:Standby supervisor failed to boot up.
```

This error message is also generated if one of the following situations apply:

- You remain at the `loader>` prompt for an extended period of time.
- You do not set the boot variables appropriately.

Replacing Modules

When you replace any module (supervisor, switching, or services module), you must ensure that the new module is running the same software version as the rest of the switch.

Refer to *Cisco MDS 9000 Family SAN Volume Controller Configuration Guide* for configuration details on replacing the caching services module (CSM).



Note

When a spare standby supervisor module is inserted, it uses the same image as the active supervisor module. The Cisco SAN-OS software image is not automatically copied to the standby flash device.



Tip

Issue the **install all** command to copy the Cisco SAN-OS software image to the standby supervisor bootflash device.

Issuing the **install all** command after replacing any module, ensures the following actions:

- The proper system and kickstart images are copied on the standby bootflash: file system.
- The proper boot variables are set.
- The loader and the BIOS are upgraded to the same version available on the active supervisor module.

To replace a module in any switch in the Cisco MDS 9200 Series or 9500 Series, follow these steps:

- Step 1** Create a backup of your existing configuration file, if required, using the **copy running-config startup-config** command.
- Step 2** Replace the required module as specified in the *Cisco MDS 9200 Series Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide*.

Send documentation comments to mdsfeedback-doc@cisco.com.

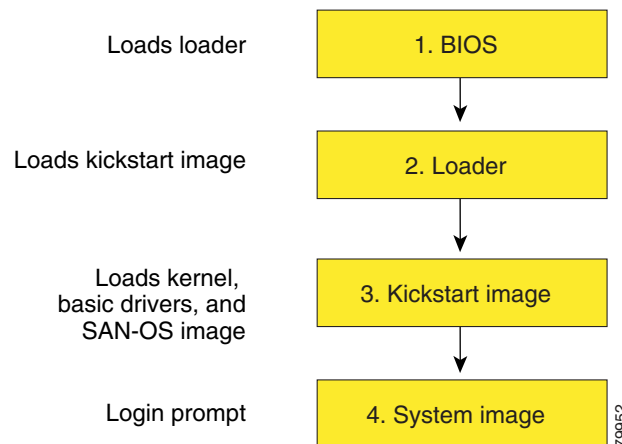
- Step 3** Verify that space is available on the standby supervisor bootflash using the **dir bootflash://sup-standby/** command. It is a good practice to remove older versions of Cisco MDS SAN-OS images and kickstart images. For information about displaying file systems and removing files, see the “Using the File System” section on page 2-22.
- Step 4** Issue the **install all** command to ensure that the new module is running the same software as the rest of the switch.
- Step 5** Wait until the new module is online and then ensure that the replacement was successful using the **show module** command.

Corrupted Bootflash Recovery

All switch configurations reside in the internal bootflash. If you have a corrupted internal bootflash, you could potentially lose your configuration. Be sure to save and back up your configuration files periodically. The regular switch boot goes through the following sequence (see Figure 6-2):

1. The basic input/output system (BIOS) loads the loader.
2. The loader loads the kickstart image into RAM and starts the kickstart image.
3. The kickstart image loads and starts the system image.
4. The system image reads the startup configuration file.

Figure 6-2 Regular Boot Sequence



If the images on your switch are corrupted and you cannot proceed (error state), you can interrupt the switch boot sequence and recover the image by entering the BIOS configuration utility described in the following section. Access this utility only when needed to recover a corrupted internal disk.



Caution

The BIOS changes explained in this section are only required to recover a corrupted bootflash.

Recovery procedures require the regular sequence to be interrupted. The internal switch sequence goes through four phases between the time you turn the switch on and the time the switch prompt appears on your terminal—BIOS, boot loader, kickstart, and system (see Table 6-2 and Figure 6-3).

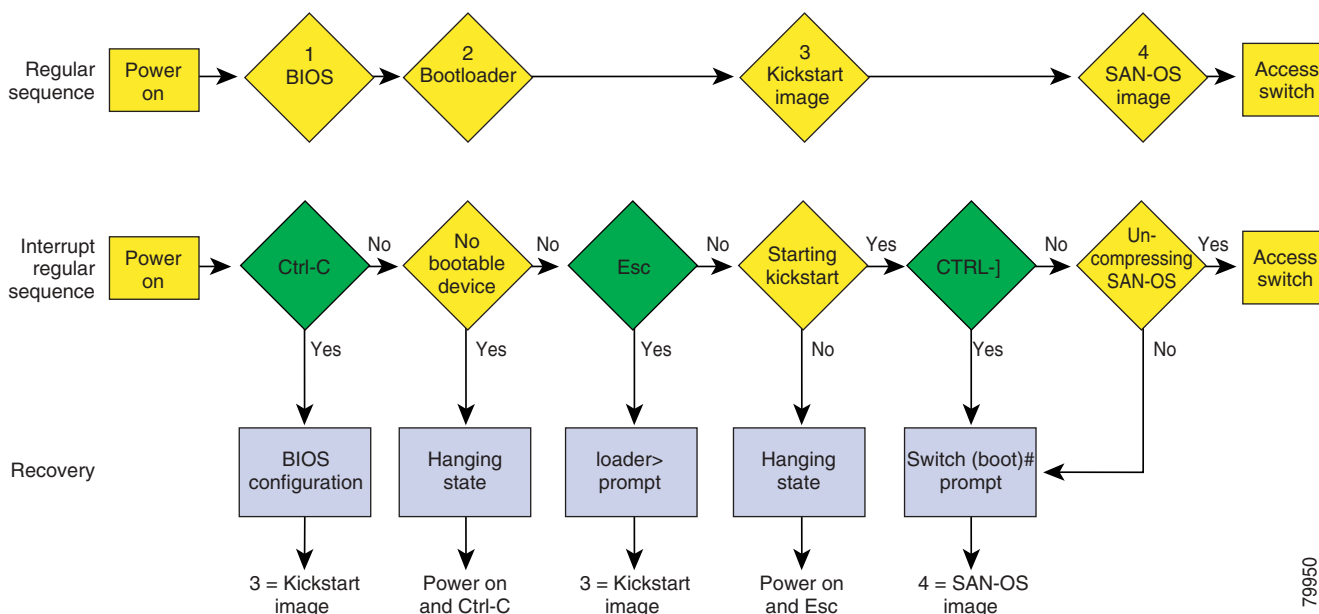
Send documentation comments to mdsfeedback-doc@cisco.com.

Table 6-2 Recovery Interruption

| Phase | Normal Prompt ¹ | Recovery Prompt ² | Description |
|-------------|----------------------------|------------------------------|---|
| BIOS | loader> | No bootable device | The BIOS begins the power-on self test, memory test, and other operating system applications. While the test is in progress, press Ctrl-C to enter the BIOS configuration utility and use the netboot option. |
| Boot loader | Starting kickstart | loader> | The boot loader uncompresses loaded software to boot an image using its file name as reference. These images are made available through bootflash. When the memory test is over, press Esc to enter the boot loader prompt. |
| Kickstart | Uncompressing system | switch (boot) # | When the boot loader phase is over, press Ctrl-] ³ (Control key plus right bracket key) to enter the <code>switch (boot) #</code> prompt. If the corruption causes the console to stop at this prompt, copy the system image and reboot the switch. |
| System | Login: | — | The system image loads the configuration file of the last saved running configuration and returns a switch login prompt. |

1. This prompt or message appears at the end of each phase.
2. This prompt or message appears when the switch cannot progress to the next phase.
3. Depending on your Telnet client, these keys may be reserved and you need to remap the keystroke. Refer to the documentation provided by your Telnet client.

Figure 6-3 Regular and Recovery Sequence



79950

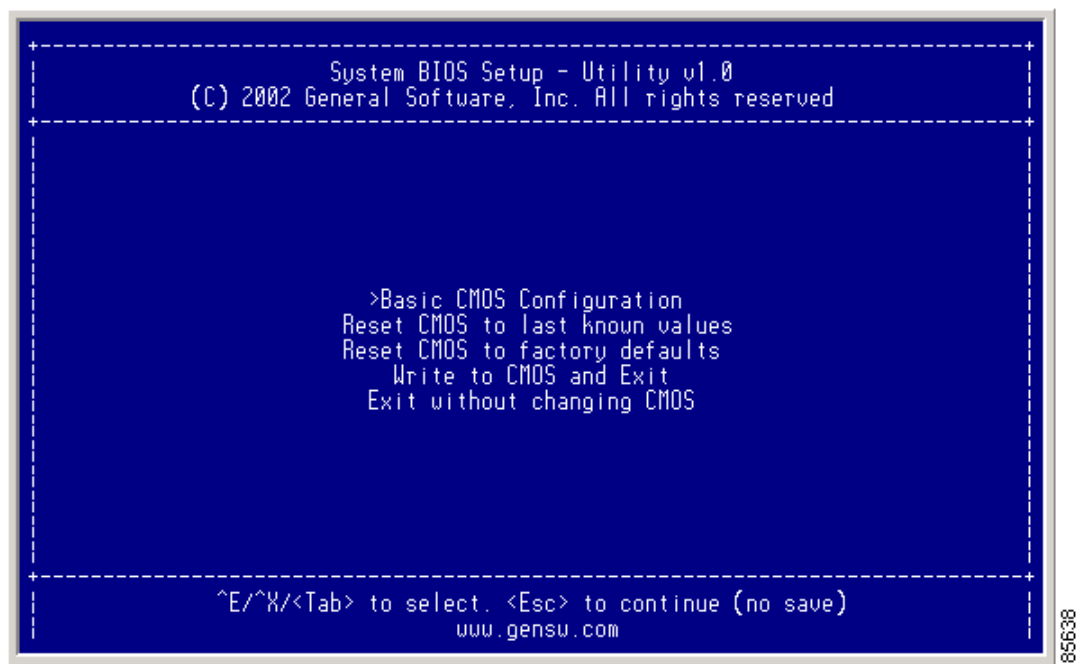
Send documentation comments to mdsfeedback-doc@cisco.com.

Recovery Using BIOS Setup

To recover a corrupted bootflash image (no bootable device found message) for a switch with a single supervisor module, follow these steps:

-
- Step 1** Connect to the console port of the required switch.
 - Step 2** Boot or reboot the switch.
 - Step 3** Press **Ctrl-C** to interrupt the BIOS setup during the BIOS memory test.
You see the netboot BIOS Setup Utility screen (see [Figure 6-4](#)).

Figure 6-4 *BIOS Setup Utility*



Note

Your navigating options are provided at the bottom of the screen.

Tab = Jump to next field

Ctrl-E = Down arrow

Ctrl-X = Up arrow

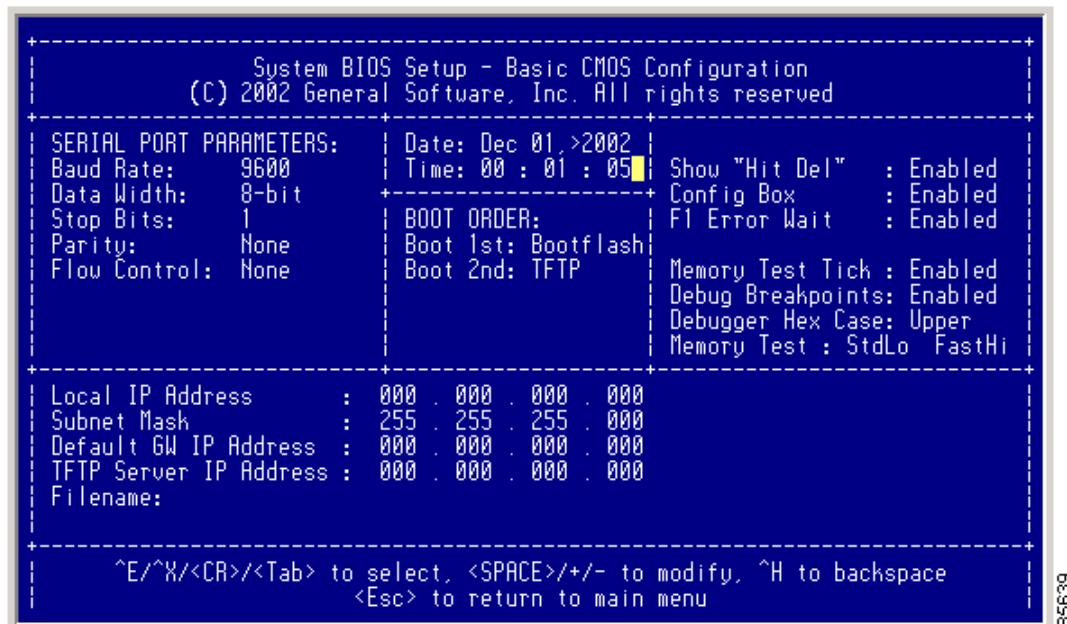
Ctrl-H = Erase (Backspace might not work if your terminal is not configured properly.)

Send documentation comments to mdsfeedback-doc@cisco.com.

Step 4 Press the **Tab** key to select the Basic CMOS Configuration, and press **Enter**.

You see the BIOS Setup CMOS Configuration screen (see Figure 6-5).

Figure 6-5 BIOS Setup Configuration (CMOS)



Step 5 Change the Boot 1st: field to **TFTP**.

Step 6 Press the **Tab** key until you reach the Local IP Address field.

Step 7 Enter the local IP address for the switch, and press the **Tab** key.

Step 8 Enter the subnet mask for the IP address, and press the **Tab** key.

Step 9 Enter the IP address of the default gateway, and press the **Tab** key.

Step 10 Enter the IP address of the TFTP server, and press the **Tab** key.

Step 11 Enter the image name (kickstart), and press the **Tab** key. This path should be relative to the TFTP server root directory.



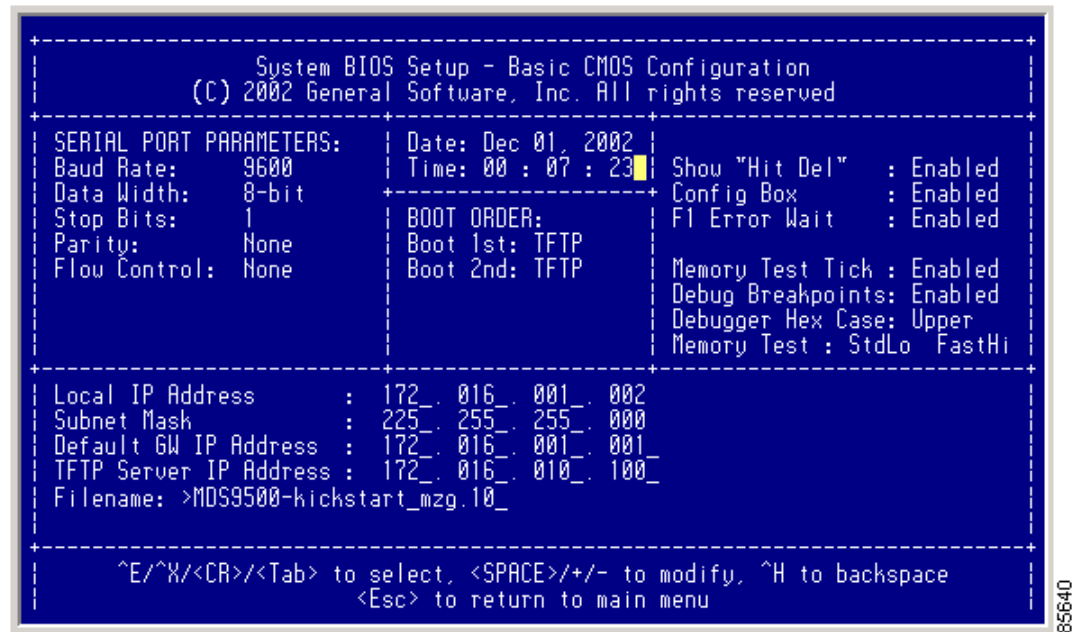
Caution

The file name must be entered exactly as it is displayed on your TFTP server. For example, if you have a file name **MDS9500-kickstart_mzg.10**, then enter this name using the exact uppercase characters and file extensions as shown on your TFTP server.

Send documentation comments to mdsfeedback-doc@cisco.com.

You see the configured changes (see Figure 6-6).

Figure 6-6 BIOS Setup Configuration (CMOS) Changes



Step 12 Press the **Esc** key to return to the main menu.

Step 13 Choose **Write to CMOS and Exit** from the main screen to save your changes.



Note These changes are saved in the CMOS.



Caution The switch must have IP connectivity to reboot using the newly configured values.

You see the following prompt:

```
switch boot) #
```

Step 14 Enter the **init system** command at the `switch boot) #` prompt, and press **Enter**.

```
switch boot) # init system
```

The `switch boot) #` prompt indicates that you have a usable kickstart image.



Note The **init system** command also installs a new loader from the existing (running) kickstart image.

Send documentation comments to mdsfeedback-doc@cisco.com.

- Step 15** Enter the **init system check-filesystem** command at the `switch(boot)#` prompt, and press **Enter**. As of Cisco MDS SAN-OS Release 2.1(1a), this command checks all the internal file systems and fixes any errors that are encountered.

```
switch(boot)# init system check-filesystem
```

- Step 16** Follow the procedure specified in the “[Recovery from the switch\(boot\)# Prompt](#)” section on page 6-31.

Recovery from the loader> Prompt

To recover a corrupted kickstart image (system error state) for a switch with a single supervisor module, follow these steps:

- Step 1** Press the **Esc** key to interrupt the boot loader setup after the BIOS memory test.



Note Press **Esc** immediately after you see the following message:

```
00000589K Low Memory Passed
00000000K Ext Memory Passed
Hit ^C if you want to run SETUP....
Wait.....
```

If you wait too long, you will skip the boot loader phase and enter the kickstart phase.

You see the `loader>` prompt.



Caution The `loader>` prompt is different from the regular `switch#` or `switch(boot)#` prompt. The CLI command completion feature does not work at this prompt and may result in undesired errors. You must type the command exactly as you want the command to appear.



Tip Use the **help** command at the `loader>` prompt to display a list of commands available at this prompt or to obtain more information about a specific command in that list.

- Step 2** Enter the local IP address and the subnet mask for the switch, and press **Enter**.

```
loader> ip address 172.16.1.2 255.255.255.0
Found Intel EtherExpressPro100 82559ER at 0xe800, ROM address 0xc000
Probing...[Intel EtherExpressPro100 82559ER]Ethernet addr: 00:05:30:00:52:27
Address: 172.16.1.2
Netmask: 255.255.255.0
Server: 0.0.0.0
Gateway: 0.0.0.0
```

- Step 3** Enter the IP address of the default gateway, and press **Enter**.

```
loader> ip default-gateway 172.16.1.1
Address: 172.16.1.2
Netmask: 255.255.255.0
Server: 0.0.0.0
Gateway: 172.16.1.1
```


Send documentation comments to mdsfeedback-doc@cisco.com.

Step 4 Boot the kickstart image file from the required server, and press **Enter**.

```
loader> boot tftp://172.16.1.2/kickstart-latest
Address: 172.16.1.2
Netmask: 255.255.255.0
Server: 172.16.10.100
Gateway: 172.16.1.1
Booting: /kick-282 console=ttyS0,9600n8nn quiet loader_ver= "1.0(2)"....
.....Image verification OK
Starting kernel...
INIT: version 2.78 booting
Checking all filesystems..... done.
Loading system software
INIT: Sending processes the TERM signal
Sending all processes the TERM signal... done.
Sending all processes the KILL signal... done.
Entering single-user mode...
INIT: Going single user
INIT: Sending processes the TERM signal
switch(boot)#
```

The `switch(boot)#` prompt indicates that you have a usable Kickstart image.

Step 5 Copy the system and kickstart images again.

```
switch(boot)# copy scp://user@172.16.10.100/system-img bootflash:system-img
Trying to connect to tftp server.....

switch(boot)# copy scp://user@172.16.10.100/kickstart-img bootflash:kickstart-img
Trying to connect to tftp server.....
```

Step 6 Follow the procedure specified in the “[Recovery from the switch\(boot\)# Prompt](#)” section on page 6-31.

Recovery from the switch(boot)# Prompt

To recover a system image using the kickstart image for a switch with a single supervisor module, follow these steps:

Step 1 Follow this step if you issued an **init system** command. Otherwise, skip to [Step 2](#).

- a. Change to configuration mode and configure the IP address of the switch’s `mgmt0` interface.

```
switch(boot)# config t
switch(boot) (config)# interface mgmt0
```

- b. Enter the local IP address and the subnet mask for the switch, and press **Enter**.

```
switch(boot) (config-mgmt0)# ip address 172.16.1.2 255.255.255.0
```

Step 2 Issue the **no shut** command to enable the interface on the switch, and press **Enter**.

```
switch(boot) (config-mgmt0)# no shut
```

Step 3 Follow this step if you issued an **init system** command. Otherwise, skip to [Step 4](#).

- a. Enter the IP address of the default gateway, and press **Enter**.

```
switch(boot) (config-mgmt0)# ip default-gateway 172.16.1.1
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Step 4 Exit to EXEC mode.

```
switch(boot) (config-mgmt0) # end
```

Step 5 Enter the **init system check-filesystem** command at the `switch(boot)#` prompt, and press **Enter**. As of Cisco MDS SAN-OS Release 2.1(1a), this command checks all the internal file systems and fixes any errors that are encountered.

```
switch(boot) # init system check-filesystem
```

Step 6 Copy the system image from the required TFTP server, and press **Enter**.

```
switch(boot) # copy scp://user@172.16.10.100/system-img bootflash:system-img
```

Step 7 Copy the kickstart image from the required TFTP server, and press **Enter**.

```
switch(boot) # copy scp://user@172.16.10.100/kickstart-img bootflash:kickstart-img
```

Step 8 Verify that the system and kickstart image files are copied to your bootflash: filesystem.

```
switch(boot) # dir bootflash:
40295206   Aug 05 15:23:51 1980   ilc1.bin
12456448   Jul 30 23:05:28 1980   kickstart-image1
12288      Jun 23 14:58:44 1980   lost+found/
27602159   Jul 30 23:05:16 1980   system-image1
12447232   Aug 05 15:08:30 1980   kickstart-image2
28364853   Aug 05 15:11:57 1980   system-image2
```

```
Usage for bootflash://sup-local
135404544 bytes used
49155072 bytes free
184559616 bytes total
```

Step 9 Load the system image from the bootflash: filesystem.

```
switch(boot) # load bootflash:system-image
Uncompressing system image: bootflash:/system-image
CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
```

```
Would you like to enter the initial configuration mode? (yes/no): yes
```

See the “[Initial Setup Routine](#)” section on page 4-2.



Note If you enter **no** at this point, you will return to the `switch#` login prompt, and you must manually configure the switch.

Recovery for Switches with Dual Supervisor Modules

If one supervisor module is functioning and the other is not, boot the functioning supervisor module. Then use the booted supervisor module to bring up the supervisor module that is stuck. At the switch prompt, issue the **reload module slot force-dnld** command, where *slot* is the slot number of the stuck supervisor module.

If both supervisor modules are not functioning, treat it like a single supervisor module recovery. First recover the image on one supervisor module and then follow the single supervisor recovery process.

Send documentation comments to mdsfeedback-doc@cisco.com.



Note

If you do not issue the **reload module** command when a boot failure has occurred, the active supervisor module automatically reloads the standby supervisor module within 3 to 6 minutes after the failure (see the “Standby Supervisor Boot Alert” section on page 6-24).

Recognizing Error States

If you see the error messages displayed in [Figure 6-7](#) or [Figure 6-8](#), follow the procedure specified in the “Recovery Using BIOS Setup” section on page 6-27.

Figure 6-7 Error State if Powered On and Ctrl-C Is Entered

```

+-----+
|          System BIOS Configuration, (C) 2002 General Software, Inc.          |
+-----+
| System CPU       : Pentium III | Low Memory       : 630KB |
| Coprocessor     : Enabled    | Extended Memory  : 957MB |
| Embedded BIOS Date : 09/10/02 | ROM Shadowing   : Enabled |
+-----+
Boot network name is EOBC
Local IP address: 127.1.2.1

Bind to network device '/DEV/TCPIP/EOBC/BootNet'
SoBindNetName: KeOpenFile failed.
Cannot bind to the network '/DEV/TCPIP/EOBC/BootNet'
Could not get BOOTP response from the server.
BOOTNET: Dispatch duration could not be restored, reason=1.
Network boot failed, status=317.

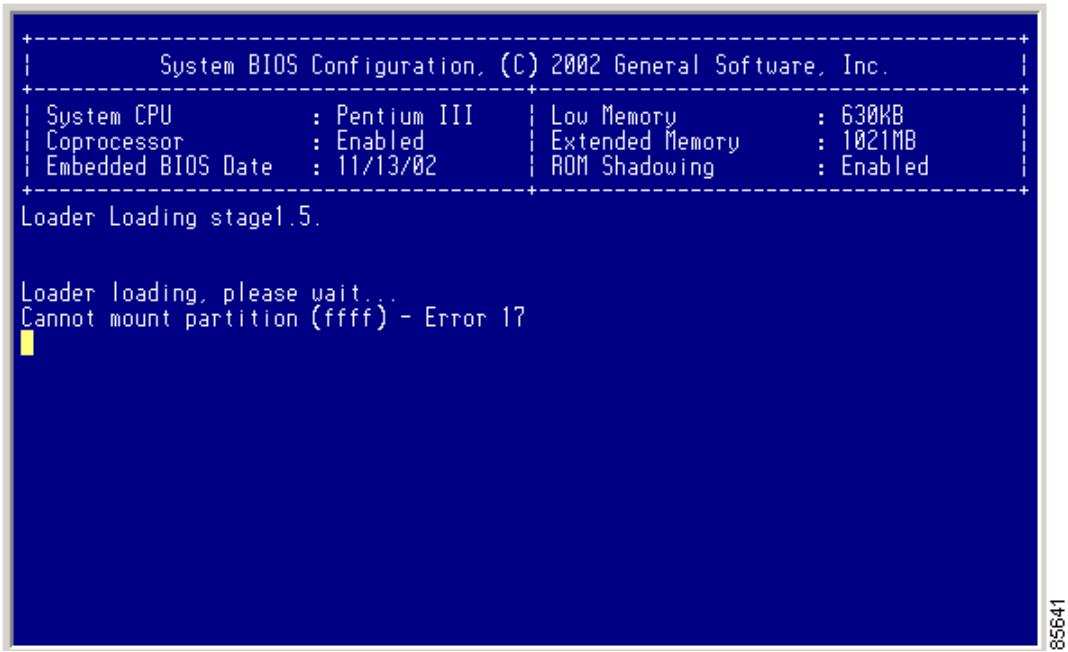
No bootable device available.
R - REBOOT
S - SETUP
ESC - BIOS DEBUGGER

```

85642

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 6-8 Error State if Powered On and Esc Is Pressed



Default Settings

Table 6-3 lists the default image settings for all Cisco MDS 9000 Family switches.

Table 6-3 Default Image Settings

| Parameters | Default |
|-----------------|------------------------|
| Kickstart image | No image is specified. |
| System image | No image is specified. |



Managing Modules

This chapter describes how to manage switching and services modules (also known as line cards) and provides information on monitoring module states. This chapter includes the following sections:

- [About Modules, page 7-1](#)
- [Verifying the Status of a Module, page 7-3](#)
- [Checking the State of a Module, page 7-4](#)
- [Connecting to a Module, page 7-4](#)
- [Reloading Modules, page 7-6](#)
- [Preserving Module Configuration, page 7-7](#)
- [Purging Module Configuration, page 7-8](#)
- [Powering Off Switching Modules, page 7-8](#)
- [Identifying Module LEDs, page 7-9](#)
- [EPLD Configuration, page 7-13](#)
- [Installing the ASM and Specifying the ASM Image Boot Variable, page 7-15](#)
- [Installing the SSM and Specifying the SSI Image Boot Variable, page 7-18](#)
- [Default Settings, page 7-23](#)

About Modules

[Table 7-1](#) describes the supervisor module options for switches in the Cisco MDS 9000 Family.

Table 7-1 **Supervisor Module Options**

| Product | No. of Supervisor Modules | Supervisor Module Slot No. | Switching and Services Module Features |
|-----------------|---------------------------|----------------------------|--|
| Cisco MDS 9216 | One module | 1 | 2-slot chassis allows one optional switching or services module in the other slot. |
| Cisco MDS 9216A | One module | 1 | 2-slot chassis allows one optional switching or services module in the other slot. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 7-1 Supervisor Module Options (continued)

| Product | No. of Supervisor Modules | Supervisor Module Slot No. | Switching and Services Module Features |
|-----------------|---------------------------|----------------------------|--|
| Cisco MDS 9216i | One module | 1 | 2-slot chassis allows one optional switching or services module in the other slot. |
| Cisco MDS 9509 | Two modules | 5 and 6 | 9-slot chassis allows any switching or services module in the other seven slots. |
| Cisco MDS 9506 | Two modules | 5 and 6 | 6-slot chassis allows any switching or services module in the other four slots. |

Supervisor Modules

Supervisor modules are automatically powered up and started with the switch.

- Cisco MDS 9200 Series switches have one supervisor module that includes an integrated 16-port switching module.
- Cisco MDS 9216i Switches have one supervisor module that includes an integrated switching module with 14 Fibre Channel ports and two Gigabit Ethernet ports.
- Cisco MDS 9500 Series switches have two supervisor modules—one in slot 5 (sup-1) and one in slot 6 (sup-2). See [Table 7-2](#). When the switch powers up and both supervisor modules come up together, the active module is the one that comes up first. The standby module constantly monitors the active module. If the active module fails, the standby module takes over without any impact to user traffic.

Table 7-2 Supervisor Module Terms and Usage in Console Displays

| Module Terms | Fixed or Relative | Usage |
|----------------------------|---------------------------------------|---|
| module-5 and module-6 | fixed usage for MDS 9509 and MDS 9506 | module-5 always refers to the supervisor module in slot 5 and module-6 always refers to the supervisor module in slot 6. |
| module-1 | fixed usage for MDS 9200 series | module-1 always refers to the supervisor module in slot 1. |
| sup-1 and sup-2 | fixed usage | sup-1 always refers to the supervisor module in slot 5 and sup-2 always refers to the supervisor module in slot 6. |
| sup-active and sup-standby | relative usage | sup-active refers to the active supervisor module—relative to the slot that contains the active supervisor module. sup-standby refers to the standby supervisor module—relative to the slot that contains the standby supervisor module. |
| sup-local and sup-remote | relative usage | If you are logged into the active supervisor, sup-local refers to the active supervisor module and sup-remote refers to the standby supervisor module. If you are logged into the standby supervisor, sup-local refers to the standby supervisor module (the one you are logged into.) There is no sup-remote available from the standby supervisor module (you cannot access a filesystem on the active sup). |

Send documentation comments to mdsfeedback-doc@cisco.com.

Switching Modules

Cisco MDS 9000 Family switches support any switching module in any non-supervisor slot. These modules obtain their image from the supervisor module.

Services Modules

Cisco MDS 9000 Family switches support any services module in any non-supervisor slot.

Refer to the *Cisco MDS 9000 Family SAN Volume Controller Configuration Guide* for more information on CSMs.

Verifying the Status of a Module

Before you begin configuring the switch, you need to ensure that the modules in the chassis are functioning as designed. To verify the status of a module at any time, issue the **show module** command (see the “[Fibre Channel Interfaces](#)” section on page 12-2). The interfaces in each module are ready to be configured when the `ok` status is displayed in the **show module** command output. A sample output of the **show module** command follows:

```
switch# show module
Mod  Ports  Module-Type                Model                Status
---  -
2    8      IP Storage Services Module DS-X9308-SMIP        ok
4    0      Caching Services Module   DS-X9530-SF1-K9      ok
5    0      Supervisor/Fabric-1       DS-X9530-SF1-K9      active *
6    0      Supervisor/Fabric-1       DS-X9530-SF1-K9      ha-standby
8    0      Caching Services Module   DS-X9560-SMAP        ok
9    32     1/2 Gbps FC Module        DS-X9032              ok
```

```
Mod  Sw          Hw          World-Wide-Name(s) (WWN)
---  -
2    1.3(0.106a) 0.206       20:41:00:05:30:00:00:00 to 20:48:00:05:30:00:00:00
5    1.3(0.106a) 0.602       --
6    1.3(0.106a) 0.602       -- <----- New running version in module 6
8    1.3(0.106a) 0.702       --
9    1.3(0.106a) 0.3         22:01:00:05:30:00:00:00 to 22:20:00:05:30:00:00:00
```

```
Mod  MAC-Address(es)                Serial-Num
---  -
2    00-05-30-00-9d-d2 to 00-05-30-00-9d-de JAB064605a2
5    00-05-30-00-64-be to 00-05-30-00-64-c2
6    00-d0-97-38-b3-f9 to 00-d0-97-38-b3-fd JAB06350B1R
8    00-05-30-01-37-7a to 00-05-30-01-37-fe JAB072705ja
9    00-05-30-00-2d-e2 to 00-05-30-00-2d-e6 JAB06280ae9
```

* this terminal session

The Status column in the output should display an `ok` status for switching modules and an active or standby (or HA-standby) status for supervisor modules. If the status is either `ok` or `active`, you can continue with your configuration.

Send documentation comments to mdsfeedback-doc@cisco.com.


Note

A standby supervisor module reflects the HA-standby status if the HA switchover mechanism is enabled (see the [“HA Switchover Characteristics” section on page 5-2](#)). If the warm switchover mechanism is enabled, the standby supervisor module reflects the standby status.

The states through which a switching module progresses is discussed in the [“Checking the State of a Module” section on page 7-4](#).

Checking the State of a Module

If your chassis has more than one switching module (also known as line card), you can check the progress by issuing the **show module** command several times and viewing the `status` column each time.

The switching module goes through a testing and an initializing stage before displaying an `ok` status. [Table 7-3](#) describes the possible states in which a module can exist.

Table 7-3 Module States

| show module Command Status Output | Description |
|--|---|
| powered up | The hardware has electrical power. When the hardware is powered up, the software begins booting. |
| testing | The switching module has established connection with the supervisor supervisor and the switching module is performing bootup diagnostics. |
| initializing | The diagnostics have completed successfully and the configuration is being downloaded. |
| failure | The switch detects a switching module failure upon initialization and automatically attempts to power-cycle the module three times. After the third attempt it continues to display a failed state. |
| ok | The switch is ready to be configured. |
| power-denied | The switch detects insufficient power for a switching module to power up (see the “Displaying Environment Information” section on page 8-11). |
| active | This module is the active supervisor module and the switch is ready to be configured. |
| HA-standby | The HA switchover mechanism is enabled on the standby supervisor module (see the “HA Switchover Characteristics” section on page 5-2). |
| standby | The warm switchover mechanism is enabled on the standby supervisor module (see the “HA Switchover Characteristics” section on page 5-2). |

Connecting to a Module

At any time, you can connect to any module using the **attach module** command. Once you are at the module prompt, you can obtain further details about the module using module-specific commands in EXEC mode.

Send documentation comments to mdsfeedback-doc@cisco.com.

To attach to a module, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# attach module 6 switch(standby)# | Provides direct access to the specified module (in this example, the standby supervisor module is in slot 6). |
| Step 2 | switch(standby)# dir bootflash: <pre> root 14502912 Jan 13 12:23:52 1980 kickstart_image1 admin 14424576 Jan 14 06:47:29 1980 kickstart_image2 admin 14469632 Jan 14 01:29:16 1980 kickstart_image3 root 14490112 Jan 08 07:25:50 1980 kickstart_image4 root 12288 Jan 16 15:49:24 1980 lost+found/ admin 14466048 Jan 14 02:40:16 1980 kickstart_image5 admin 24206675 Jan 14 02:57:03 1980 m9500-sf1ek.bin root 19084510 Jan 13 12:23:28 1980 system_image1 admin 19066505 Jan 14 06:45:16 1980 system_image2 admin 18960567 Jan 14 01:25:21 1980 system_image5 Usage for bootflash: filesystem 158516224 bytes total used 102400 bytes free 167255040 bytes available </pre> | Provides the available space information for the standby supervisor module. Note Type exit to exit the module-specific prompt. Tip If you are not accessing the switch from a console terminal, this is the only way to access the standby supervisor module. |

You can also use the **attach module** command as follows:

- To display the standby supervisor module information, although you cannot configure the standby supervisor module using this command.
- To display the switching module portion of the Cisco MDS 9200 Series supervisor module which resides in slot 1.

Send documentation comments to mdsfeedback-doc@cisco.com.

Reloading Modules

You can reload the entire switch, reset specific modules in the switch, or reload the image on specific modules in the switch.

Reloading the Switch

To reload the switch, issue the **reload** command without any options. When you issue this command, you reboot the switch (see [Chapter 6, “Software Images”](#)).



Note

If you need to issue the **reload** command, be sure to save the running configuration using the **copy running-config startup-config** command.

Power Cycling Modules

To power cycle any module, follow these steps:

-
- Step 1** Identify the module that needs to be reset.
 - Step 2** Issue the **reload module** command to reset the identified module. This command merely power cycles the selected module.

```
switch# reload module number
```

Where *number* indicates the slot in which the identified module resides. For example:

```
switch# reload module 2
```

Reloading Switching Modules

Switching modules automatically download their images from the supervisor module and do not need a forced download. This procedure is provided for reference should a need arise.

To replace the image on a switching module, follow these steps:

-
- Step 1** Identify the switching module that requires the new image.
 - Step 2** Issue the **reload module number force-dnld** command to update the image on the switching module.

```
switch# reload module number force-dnld
```

Where *number* indicates the slot in which the identified module resides. In this example, the identified module resides in slot 9.

```
switch# reload module 9 force-dnld...
Jan 1 00:00:46 switch %LC-2-MSG:SLOT9 LOG_LC-2-IMG_DNLD_COMPLETE: COMPLETED
downloading of linecard image. Download successful...
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Preserving Module Configuration

Issue the **copy running-config startup-config** command from EXEC mode to save the new configuration into nonvolatile storage. Once this command is issued, the running and the startup copies of the configuration are identical.

Table 7-4 displays various scenarios when module configurations are preserved or lost.

Table 7-4 Switching Module Configuration Status

| Scenario | Consequence |
|---|---|
| A particular switching module is removed and the copy running-config startup-config command is issued again. | The configured module information is lost. |
| A particular switching module is removed and the same switching module is replaced before the copy running-config startup-config command is issued again. | The configured module information is preserved. |
| A particular switching module is removed and replaced with the same type switching module, and a reload module number command is issued. | The configured module information is preserved. |
| A particular switching module is reloaded when a reload module number command is issued. | The configured module information is preserved. |
| A particular switching module is removed and replaced with a different type of switching module. For example, a 16-port switching module is replaced with a 32-port switching module. | The configured module information is lost from the running configuration. The default configuration is applied. The configured module information remains in startup configuration until a copy running-config startup-config command is issued again. |
| <p>Sample scenario:</p> <ol style="list-style-type: none"> 1. The switch currently has a 16-port switching module and the startup and running configuration files are the same. 2. You replace the 16-port switching module in the switch with a 32-port switching module. 3. Next, you remove the 32-port switching module and replace it with the same 16-port switching module referred to in Step 1. 4. You reload the switch. | <p>Sample response:</p> <ol style="list-style-type: none"> 1. The switch uses the 16-port switching module and the present configuration is saved in nonvolatile storage. 2. The factory default configuration is applied. 3. The factory default configuration is applied. 4. The configuration saved in nonvolatile storage referred to in Step 1 is applied. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Purging Module Configuration

Issue the **purge module *slot* running-config** command from EXEC mode to delete the configuration in a specific module. Once this command is issued, the running configuration is cleared for the specified slot. This command does not work on supervisor modules or on any slot that currently has a module. This command only works on an empty slot (where the specified module once resided).

The **purge module** command clears the configuration for any module that previously existed in a slot and has since been removed. While the module was in that slot, some parts of the configuration may have been stored in the running configuration and cannot be reused (for example, IP addresses), unless it is cleared from the running configuration.

For example, suppose you create an IP storage configuration with an IPS module in slot 3 in Switch A. This module uses IP address 10.1.5.500. You decide to remove this IPS module and move it to Switch B, and you no longer need the IP address 10.1.5.500. If you try to configure this unused IP address, you will receive an error message that prevents you from proceeding with the configuration. In this case, you need to issue the **purge module 3 running-config** command to clear the old configuration in Switch A before proceeding with using this IP address.

Powering Off Switching Modules

By default, all switching modules are configured to be in the power up state.

To power off a module, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# poweroff module 1 switch(config)# | Powers off the specified module (switching module 1) in the switch. |
| | switch(config)# no poweroff module 1 switch(config)# | Powers up the specified module (switching module 1) in the switch. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Identifying Module LEDs

Table 7-5 describes the LEDs for the Cisco MDS 9200 Series integrated supervisor modules.

Table 7-5 LEDs for the Cisco MDS 9200 Series Supervisor Modules

| LED | Status | Description |
|--------|-----------------|---|
| Status | Green | All diagnostics pass. The module is operational (normal initialization sequence). |
| | Orange | The module is booting or running diagnostics (normal initialization sequence). or The inlet air temperature of the system has exceeded the maximum system operating temperature limit (a minor environmental warning). To ensure maximum product life, you should immediately correct the environmental temperature and restore the system to normal operation. |
| | Red | The diagnostic test failed. The module is not operational because a fault occurred during the initialization sequence. or The inlet air temperature of the system has exceeded the safe operating temperature limits of the card (a major environmental warning). The card has been shut down to prevent permanent damage. The system will be shut down after two minutes if this condition is not cleared. |
| Speed | On | 2-Gbps mode and beacon mode disabled. |
| | Off | 1-Gbps mode and beacon mode disabled. |
| | Flashing | Beacon mode enabled. See the “Identifying the Beacon LEDs” section on page 12-17 . |
| Link | Solid green | Link is up. |
| | Solid yellow | Link is disabled by software. |
| | Flashing yellow | A fault condition exists. |
| | Off | No link. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 7-6 describes the LEDs for the Cisco MDS 9200 Series interface module.

Table 7-6 LEDs on the Cisco MDS 9200 Series Interface Module

| LED | Status | Description |
|-----------------------------------|--------|---|
| Status | Green | All diagnostics pass. The module is operational (normal initialization sequence). |
| | Orange | The module is booting or running diagnostics (normal initialization sequence). or The inlet air temperature of the system has exceeded the maximum system operating temperature limit (a minor environmental warning). To ensure maximum product life, you should immediately correct the environmental temperature and restore the system to normal operation. |
| | Red | The diagnostic test failed. The module is not operational because a fault occurred during the initialization sequence. or The inlet air temperature of the system has exceeded the safe operating temperature limits of the card (a major environmental warning). The card has been shut down to prevent permanent damage. |
| System | Green | All chassis environmental monitors are reporting OK. |
| | Orange | The power supply failed or the power supply fan failed. or Incompatible power supplies are installed. or The redundant clock failed. |
| | Red | The temperature of the supervisor module exceeded the major threshold. |
| MGMT 10/100 Ethernet Link LED | Green | Link is up. |
| | Off | No link. |
| MGMT 10/100 Ethernet Activity LED | Green | Traffic is flowing through port. |
| | Off | No link or no traffic. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 7-7 describes the LEDs for the 16-port and 32-port switching modules.

Table 7-7 LEDs for the Cisco MDS 9000 Family Fibre Channel Switching Modules

| LED | Status | Description |
|--------|-----------------------------|---|
| Status | Green | All diagnostics pass. The module is operational (normal initialization sequence). |
| | Orange | The module is booting or running diagnostics (normal initialization sequence). or The inlet air temperature of the system has exceeded the maximum system operating temperature limit (a minor environmental warning). To ensure maximum product life, you should immediately correct the environmental temperature and restore the system to normal operation. |
| | Red | The diagnostic test failed. The module is not operational because a fault occurred during the initialization sequence. or The inlet air temperature of the system has exceeded the safe operating temperature limits of the card (a major environmental warning). The card has been shut down to prevent permanent damage. |
| Speed | On | 2-Gbps mode. |
| | Off | 1-Gbps mode. |
| Link | Solid green | Link is up. |
| | Steady flashing green | Link is up (beacon used to identify port). |
| | Intermittent flashing green | Link is up (traffic on port). |
| | Solid yellow | Link is disabled by software. |
| | Flashing yellow | A fault condition exists. |
| | Off | No link. |

Send documentation comments to mdsfeedback-doc@cisco.com.

The LEDs on the supervisor module indicate the status of the supervisor module, power supplies, and the fan module. [Table 7-8](#) provides more information about these LEDs.

Table 7-8 LEDs for the Cisco MDS 9500 Series Supervisor Modules

| LED | Status | Description |
|-----------------------------------|--------|--|
| Status | Green | All diagnostics pass. The module is operational (normal initialization sequence). |
| | Orange | The module is booting or running diagnostics (normal initialization sequence). or An over temperature condition has occurred (a minor threshold has been exceeded during environmental monitoring). |
| | Red | The diagnostic test failed. The module is not operational because a fault occurred during the initialization sequence. or An over temperature condition occurred (a major threshold was exceeded during environmental monitoring). |
| System ¹ | Green | All chassis environmental monitors are reporting OK. |
| | Orange | The power supply has failed or the power supply fan has failed. or Incompatible power supplies are installed. or The redundant clock has failed. |
| | Red | The temperature of the supervisor module major threshold has been exceeded. |
| Active | Green | The supervisor module is operational and active. |
| | Orange | The supervisor module is in standby mode. |
| Pwr Mgmt ¹ | Green | Sufficient power is available for all modules. |
| | Orange | Sufficient power is not available for all modules. |
| MGMT 10/100 Ethernet Link LED | Green | Link is up. |
| | Off | No link. |
| MGMT 10/100 Ethernet Activity LED | Green | Traffic is flowing through port. |
| | Off | No link or no traffic. |
| CompactFlash | Green | The external CompactFlash card is being accessed. |
| | Off | No activity. |

1. The System and Pwr Mgmt LEDs on a redundant supervisor module are synchronized to the active supervisor module.

Send documentation comments to mdsfeedback-doc@cisco.com.

EPLD Configuration

Switches and directors in the Cisco MDS 9000 Family contain several electrical programmable logical devices (EPLDs) that provide hardware functionalities in all modules. Starting with Cisco MDS SAN-OS Release 1.2, EPLD image upgrades are periodically provided to include enhanced hardware functionality or to resolve known issues.



Tip

Refer to the *Cisco MDS SAN-OS Release Notes* to verify if the EPLD has changed for the Cisco SAN-OS image version being used.

EPLDs can be upgraded or downgraded using CLI commands. When EPLDs are being upgraded or downgraded, the following guidelines and observations apply:

- You can individually update each module that is online. The EPLD update is only disruptive to the module being upgraded.
- If you interrupt an upgrade, the module must be upgraded again.
- The upgrade or downgrade can only be executed from the active supervisor module. While the active supervisor module cannot be updated, you can update the other modules individually.
- In Cisco MDS 9100 Series fabric switches, be sure to specify one (1) as the module number.
- Cisco MDS 9200 Series switches do not support EPLD upgrades.



Caution

Do not insert or remove any modules while an EPLD upgrade or downgrade is in progress.

Upgrading EPLD Images

Use the **install module *number* *epld* *url*** command on the active supervisor module to upgrade EPLD images for a module.

```
switch# install module 2 epld bootflash:m9000-epld-2.0.1b.img
```

| EPLD | Curr Ver | New Ver |
|--------------------|----------|---------|
| Power Manager | 0x07 | 0x08 |
| XBUS IO | 0x03 | 0x03 |
| UD Flow Control | 0x05 | 0x05 |
| PCI ASIC I/F | 0x05 | 0x05 |
| Service Module I/F | 0x1a | 0x1a |

Module 2 will be powered down now!!

Do you want to continue (y/n) ? **y**

\ <-----progress twirl

Module 2 EPLD upgrade is successful

Send documentation comments to mdsfeedback-doc@cisco.com.

If you forcefully upgrade a module that is not online, all EPLDs are forcefully upgraded. If the module is not present in the switch, an error is returned. If the module is present, the command process continues. To upgrade a module that is not online but is present in the chassis, use the same command. The switch software prompts you to continue after reporting the module state. When you confirm your intention to continue, the upgrade continues.

```
switch# install module 2 epld bootflash:m9000-epld-2.0.1b.img
\ <-----progress twirl
Module 2 EPLD upgrade is successful
```


Note

When you upgrade the EPLD module on Cisco MDS 9100 Series switches, you receive the following message:

```
Data traffic on the switch will stop now!!
Do you want to continue (y/n) ?
```


Note

The same procedure used to upgrade the EPLD images on a module can be used to downgrade the EPLD images.

Displaying EPLD Versions

Use the **show version module number epld** command to view all current EPLD versions on a specified module (see [Example 7-1](#)).

Example 7-1 Displays Current EPLD Versions for a Specified Module

```
switch# show version module 2 epld
EPLD Device                      Version
-----
Power Manager                    0x07
XBUS IO                          0x03
UD Flow Control                  0x05
PCI ASIC I/F                     0x05
Service Module I/F               0x1a
```

Use the **show version epld url** command to view the available EPLD versions (see [Example 7-2](#)).

Example 7-2 Displays Available EPLD Versions

```
switch# show version epld bootflash:m9000-epld-2.0.1b.img
MDS series EPLD image, built on Mon Sep 20 16:39:36 2004
Module Type                      EPLD Device                      Version
-----
MDS 9500 Supervisor 1            XBUS 1 IO                        0x09
                                XBUS 2 IO                        0x0c
                                UD Flow Control                  0x05
                                PCI ASIC I/F                     0x04
1/2 Gbps FC Module (16 Port)     XBUS IO                          0x07
                                UD Flow Control                  0x05
                                PCI ASIC I/F                     0x05
1/2 Gbps FC Module (32 Port)     XBUS IO                          0x07
                                UD Flow Control                  0x05
                                PCI ASIC I/F                     0x05
Advanced Services Module         XBUS IO                          0x07
```

Send documentation comments to mdsfeedback-doc@cisco.com.

| | | |
|-------------------------------------|--------------------|------------|
| | UD Flow Control | 0x05 |
| | PCI ASIC I/F | 0x05 |
| | PCI Bridge | 0x05 |
| IP Storage Services Module (8 Port) | Power Manager | 0x07 |
| | XBUS IO | 0x03 |
| | UD Flow Control | 0x05 |
| | PCI ASIC I/F | 0x05 |
| | Service Module I/F | 0x0a |
| | IPS DB I/F | 0x1a |
| IP Storage Services Module (4 Port) | Power Manager | 0x07 |
| | XBUS IO | 0x03 |
| | UD Flow Control | 0x05 |
| | PCI ASIC I/F | 0x05 |
| | Service Module I/F | 0x1a |
| Caching Services Module | Power Manager | 0x08 |
| | XBUS IO | 0x03 |
| | UD Flow Control | 0x05 |
| | PCI ASIC I/F | 0x05 |
| | Service Module I/F | 0x72 |
| | Memory Decoder 0 | 0x02 |
| | Memory Decoder 1 | 0x02 |
| MDS 9100 Series Fabric Switch | XBUS IO | 0x03 |
| | PCI ASIC I/F | 0x40000003 |
| 2x1GE IPS, 14x1/2Gbps FC Module | Power Manager | 0x07 |
| | XBUS IO | 0x05 |
| | UD Flow Control | 0x05 |
| | PCI ASIC I/F | 0x07 |
| | IPS DB I/F | 0x1a |

Installing the ASM and Specifying the ASM Image Boot Variable

The virtualization image for an Advanced Services Module (ASM) can be specified using the ASM-SFN image boot variable for VERITAS Storage Foundation for Networks (VSNF) or, as of Cisco MDS SAN-OS Release 2.0(2b), the SSI variable for normal Fibre Channel switching.



Note

You cannot configure the ASM-SFN image boot variable and the SSI image boot variable for an ASM concurrently. The ASM can run only one or the other.

Configuring the ASM-SFN Image Boot Variable for VSNF

To configure the ASM-SFN image boot variable for VSNF, follow these steps:

- Step 1** Log into the switch through the console port, an SSH session, or a Telnet session.
- Step 2** Issue the **dir bootflash:** command to verify that the ASM-SFN software image file corresponding to your Cisco MDS SAN-OS release is present on the active supervisor module. For example, if your switch is running Cisco MDS SAN-OS Release 2.1(1a), you must have m9000-ek9-asm-sfn-mz.2.1.1a.bin in the bootflash: filesystem on the active supervisor module. You can find the ASM-SFN images at the following URL:
<http://www.cisco.com/cgi-bin/tablebuild.pl/mds9000-asm-3des>
- Step 3** If necessary, download the appropriate ASM-SFN software image file to your FTP server and copy it from an FTP server to the bootflash: file in the active supervisor module.

Send documentation comments to mdsfeedback-doc@cisco.com.

**Note**

Ensure that there is enough space available on the active and standby supervisor bootflash to store the ASM-SFN image while specifying the bootvar. The system will automatically sync the ASM-SFN image to the standby supervisor.

```
switch# copy ftp://10.1.7.2/m9000-ek9-asm-sfn-mz.2.1.1a.bin
bootflash:m9000-ek9-asm-sfn-mz.2.1.1a.bin
```

Step 4 Change to configuration mode.

```
switch# config terminal
```

Step 5 Specify the ASM-SFN image to be used.

```
switch(config)# boot asm-sfn bootflash:m9000-ek9-asm-sfn-mz.2.1.1a.bin module 4
```

**Note**

You can only specify one image for the ASM-SFN image boot variable per module.

Step 6 Change to the EXEC mode.

```
switch(config)# exit
```

Step 7 Issue the **show boot** command to display the current contents of the ASM-SFN image boot variable.

```
switch# show boot
sup-1
kickstart variable = bootflash:/boot-2-0-1-9
system variable =
bootflash:/isan-2-0-1-9;bootflash:/isan-2-0-0-181b;bootflash:/isan-2-0-0-181b
sup-2
kickstart variable = bootflash:/boot-2-0-1-9
system variable =
bootflash:/isan-2-0-1-9;bootflash:/isan-2-0-0-181b;bootflash:/isan-2-0-0-181b
Module 4
asm-sfn variable = bootflash:/m9000-ek9-asm-sfn-mz.2.1.1a.bin
```

Step 8 Save the new variable configuration so the new image is used the next time you log into the switch.

```
switch# copy running-config startup-config
```

Step 9 Reload the ASM-SFN to load the new image.

```
switch# reload module 4
reloading module 4 ...
```

The **reload** command power cycles the ASM.

Step 10 Issue the **show module** command to verify the status of the ASM.

```
switch# show module
```

| MMod | Ports | Module-Type | Model | Status |
|------|-------|--------------------------|-----------------|----------|
| 4 | 32 | Advanced Services Module | DS-X9032-SMV | ok |
| 5 | 0 | Supervisor/Fabric-1 | DS-X9530-SF1-K9 | active * |

| Mod | Sw | Hw | World-Wide-Name(s) (WWN) |
|-----|---------|------|--|
| 4 | 2.1(1a) | 0.30 | 20:c1:00:05:30:00:06:de to 20:e0:00:05:30:00:06:de |
| 5 | 2.1(1a) | 4.0 | -- |

| Mod | Application Image Description | Application Image Version |
|-----|-------------------------------|---------------------------|
| 4 | Advanced Services Module | 2.1(1a) |
| 5 | Supervisor/Fabric-1 | 4.0 |

Send documentation comments to mdsfeedback-doc@cisco.com.

```

4          ASM-SFN linecard image          2.1(1a)

Mod  MAC-Address(es)          Serial-Num
---  -
4    00-05-30-00-9e-b2 to 00-05-30-00-9e-b6  JAB06480590
5    00-0c-30-da-7c-18 to 00-0c-30-da-7c-1c  JAB080507P4

```

* this terminal session



Note

The next time you reboot the switch, the saved image is used. If you do not save the configuration, the previously saved startup configuration image is used.

Configuring the SSI Image Boot Variable for Fibre Channel Switching

To configure the SSI image boot variable for Fibre Channel switching, follow these steps:

- Step 1** Log into the switch through the console port, an SSH session, or a Telnet session.
- Step 2** Issue the **dir bootflash:** command to verify that the SSI software image file corresponding to your Cisco MDS SAN-OS release is present on the active supervisor module. For example, if your switch is running Cisco MDS SAN-OS Release 2.1(1a), you must have m9000-ek9-ssi-mz.2.1.1a.bin in the bootflash: filesystem on the active supervisor module. You can find the SSI images at the following URL:
<http://www.cisco.com/cgi-bin/tablebuild.pl/mds9000-ssi-3des>
- Step 3** If necessary, download the appropriate SSI software image file to your FTP server and copy it from an FTP server to the bootflash: filesystem on the active supervisor module.



Note

Ensure that there is enough space available on the active and standby supervisor bootflash to store the SSI image while specifying the bootvar. The system will automatically sync the SSI image to the standby supervisor.

```

switch# copy ftp://10.1.7.2/m9000-ek9-ssi-mz.2.1.1a.bin
bootflash:m9000-ek9-ssi-mz.2.1.1a.bin

```

- Step 4** Change to configuration mode.

```
switch# config terminal
```

- Step 5** Specify the SSI image to be used.

```
switch(config)# boot ssi bootflash:m9000-ek9-ssi-mz.2.1.1a.bin module 4
```



Note

You can only specify one image for the SSI variable per module.

- Step 6** Change to the EXEC mode.

```
switch(config)# exit
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Step 7 Issue the **show boot** command to display the current contents of the SSI variable.

```
switch# show boot
sup-1
kickstart variable = bootflash:/boot-2-0-1-9
system variable =
bootflash:/isan-2-0-1-9;bootflash:/isan-2-0-0-181b;bootflash:/isan-2-0-0-181b
sup-2
kickstart variable = bootflash:/boot-2-0-1-9
system variable =
bootflash:/isan-2-0-1-9;bootflash:/isan-2-0-0-181b;bootflash:/isan-2-0-0-181b
Module 4
ssi variable = bootflash:/m9000-ek9-ssi-mz.2.1.1a.bin
```

Step 8 Save the new variable configuration so the new image is used the next time you log into the switch.

```
switch# copy running-config startup-config
```

Step 9 Reload the SSI to load the new image.

```
switch# reload module 4
reloading module 4 ...
```

The **reload** command power cycles the ASM.

Step 10 Issue the **show module** command to verify the status of the ASM.

```
switch# show module
MMod  Ports  Module-Type                Model                Status
---  ---
4      32      Advanced Services Module   DS-X9032-SMV        ok
5       0      Supervisor/Fabric-1        DS-X9530-SF1-K9     active *
```

```
Mod  Sw          Hw      World-Wide-Name(s) (WWN)
---  ---
4    2.1(1a)      0.30    20:c1:00:05:30:00:06:de to 20:e0:00:05:30:00:06:de
5    2.1(1a)      4.0     --
```

```
Mod      Application Image Description      Application Image Version
-----
4          SSI linecard image                2.1(1a)
```

```
Mod  MAC-Address(es)                Serial-Num
---  ---
4    00-05-30-00-9e-b2 to 00-05-30-00-9e-b6  JAB06480590
5    00-0c-30-da-7c-18 to 00-0c-30-da-7c-1c  JAB080507P4
```

* this terminal session



Note

The next time you reboot the switch, the saved image is used. If you do not save the configuration, the previously saved startup configuration image is used.

Installing the SSM and Specifying the SSI Image Boot Variable

As of Cisco SAN-OS Release 2.0(2b), the image for a Storage Services Module (SSM) can be specified using the SSI image boot variable to support Fibre Channel switching and Intelligent Storage Services (see [Chapter 25, “Configuring Intelligent Storage Services”](#)). As of Cisco SAN-OS Release 2.1(1a), the

Send documentation comments to mdsfeedback-doc@cisco.com.

image for the SSM can also be specified using the ASM-SFN image boot variable to support VERITAS Storage Foundation for Networks (VSN). Once you set the SSI image boot variable, you do not need to reset it for upgrades or downgrades to any Cisco MDS SAN-OS release that supports the SSI image. You can use the **install all** command to configure SSM modules. See [Example 6-4 on page 6-12](#).



Note

You cannot configure the ASM-SFN image boot variable and the SSI image boot variable for an SSM concurrently. The SSM can run only one or the other.



Note

If you downgrade to a Cisco MDS SAN-OS release that does not support the SSM module, you must power down the SSM module. The boot variables for the SSM module are lost.

Configuring the ASM-SFN Image Boot Variable for VSN

To configure the ASM-SFN image boot variable for VSN, follow these steps:

- Step 1** Log into the switch through the console port, an SSH session, or a Telnet session.
- Step 2** Issue the **dir bootflash:** command to verify that the ASM-SFN software image file corresponding to your Cisco MDS SAN-OS release is present on the active supervisor module. For example, if your switch is running Cisco MDS SAN-OS Release 2.1(1a), you must have m9000-ek9-asm-sfn-mz.2.1.1a.bin in the bootflash: filesystem on the active supervisor module. You can find the ASM-SFN images at the following URL:
<http://www.cisco.com/cgi-bin/tablebuild.pl/mds9000-asm-3des>
- Step 3** If necessary, download the appropriate ASM-SFN software image file to your FTP server and copy it from an FTP server to the bootflash: filesystem in the active supervisor module or slot0:.



Note

Ensure there is enough space available on the active and standby supervisor bootflash to store the ASM-SFN image while specifying the bootvar. The system will automatically sync the ASM-SFN image to standby supervisor.

```
switch# copy ftp://10.1.7.2/m9000-ek9-asm-sfn-mz.2.1.1a.bin
bootflash:m9000-ek9-asm-sfn-mz.2.1.1a.bin
```

- Step 4** Change to configuration mode.
- Step 5** Specify the ASM-SFN image to be used.

```
switch(config)# boot asm-sfn bootflash:m9000-ek9-asm-sfn-mz.2.1.1a.bin module 4
```



Note

You can only specify one image for the ASM-SFN image boot variable per module.

- Step 6** Change to the EXEC mode.

```
switch(config)# exit
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Step 7 Issue the **show boot** command to display the current contents of the ASM-SFN image boot variable.

```
switch# show boot
sup-1
kickstart variable = bootflash:/boot-2-0-1-9
system variable = bootflash:/isan-2-0-1-9
sup-2
kickstart variable = bootflash:/boot-2-0-1-9
system variable = bootflash:/isan-2-0-1-9
Module 4
asm-sfn variable = bootflash:/m9000-ek9-asm-sfn-mz.2.1.1a.bin
```

Step 8 Save the new variable configuration so the new image is used the next time the switch reboots.

```
switch# copy running-config startup-config
```



Note If you do not save this configuration, it is lost on a switch reboot, and the SSM card is kept in the power-down state. You need to reimplement this procedure to bring the SSM card back up.

Step 9 Verify that the SSM module is installed in the switch.

Step 10 Reload the SSM to load the new image.

```
switch# reload module 4
reloading module 4 ...
```

The **reload** command power cycles the SSM.

Step 11 Issue the **show module** command to verify the status of the SSM.

```
switch# show module
MMod  Ports  Module-Type                Model                Status
---  ---
4      32      Storage Services Module    DS-X9032-SMV        ok
5      0        Supervisor/Fabric-1        DS-X9530-SF1-K9     active *
```

```
Mod  Sw          Hw      World-Wide-Name(s) (WWN)
---  ---
4     2.1(1a)      0.30    20:c1:00:05:30:00:06:de to 20:e0:00:05:30:00:06:de
5     2.1(1a)      4.0     --
```

```
Mod      Application Image Description      Application Image Version
-----
4         ASM-SFN linecard image                2.1(1a)
```

```
Mod  MAC-Address(es)                Serial-Num
---  ---
4     00-05-30-00-9e-b2 to 00-05-30-00-9e-b6  JAB06480590
5     00-0c-30-da-7c-18 to 00-0c-30-da-7c-1c  JAB080507P4
```

* this terminal session



Note The next time you reboot the switch, the saved image is used. If you do not save the configuration, the previously saved startup configuration image is used.

Send documentation comments to mdsfeedback-doc@cisco.com.

Configuring the SSI Image Boot Variable for Fibre Channel Switching and Intelligent Storage Services

To configure the SSI image boot variable for Fibre Channel switching and Intelligent Storage Services, follow these steps:

- Step 1** Log into the switch through the console port, an SSH session, or a Telnet session.
- Step 2** Issue the **dir bootflash:** command to verify that the SSI software image file corresponding to your Cisco MDS SAN-OS release is present on the active supervisor module. For example, if your switch is running Cisco MDS SAN-OS Release 2.1(1a), you must have m9000-ek9-ssi-mz.2.1.1a.bin in the bootflash: filesystem on the active supervisor module. You can download the SSI images from the following URL: <http://www.cisco.com/cgi-bin/tablebuild.pl/mds9000-ssi-3des>
- Step 3** If necessary, download the appropriate SSI software image file to your FTP server and copy it from an FTP server to the bootflash: filesystem on the active supervisor module or slot0:.



Note Ensure there is enough space available on the active and standby supervisor bootflash to store the SSI image while specifying the bootvar. The system will automatically sync the SSI image to standby supervisor.



Note You cannot set boot variables for images in slot0:.

```
switch# copy ftp://10.1.7.2/m9000-ek9-ssi-mz.2.1.1a.bin
bootflash:m9000-ek9-ssi-mz.2.1.1a.bin
```

- Step 4** Change to configuration mode.

```
switch# config terminal
```

- Step 5** Specify the SSI image to be used.

```
switch(config)# boot ssi bootflash:m9000-ek9-ssi-mz.2.1.1a.bin module 4
```



Note You can only specify one image for the SSI variable per module.

- Step 6** Change to the EXEC mode.

```
switch(config)# exit
```

- Step 7** Issue the **show boot** command to display the current contents of the SSI variable.

```
switch# show boot
sup-1
kickstart variable = bootflash:/boot-2-0-1-9
system variable = bootflash:/isan-2-0-1-9
sup-2
kickstart variable = bootflash:/boot-2-0-1-9
system variable = bootflash:/isan-2-0-1-9
Module 4
ssi variable = bootflash:/m9000-ek9-ssi-mz.2.1.1a.bin
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Step 8 Save the new variable configuration so the new image is used the next time the switch reboots.

```
switch# copy running-config startup-config
```



Note If you do not save this configuration, it is lost on a switch reboot, and the SSM card is kept in the power-down state. You need to reimplement this procedure to bring the SSM card back up.

Step 9 Verify that the SSM module is installed in the switch.

Step 10 Reload the SSM to load the new image.

```
switch# reload module 4
reloading module 4 ...
```

The **reload** command power cycles the SSM.

Step 11 Issue the **show module** command to verify the status of the SSM.

```
switch# show module
```

| Mod | Ports | Module-Type | Model | Status |
|-----|-------|-------------------------|-----------------|----------|
| 4 | 32 | Storage Services Module | DS-X9032-SSM | ok |
| 5 | 0 | Supervisor/Fabric-1 | DS-X9530-SF1-K9 | active * |

| Mod | Sw | Hw | World-Wide-Name(s) (WWN) |
|-----|---------|------|--|
| 4 | 2.1(1a) | 0.30 | 20:c1:00:05:30:00:06:de to 20:e0:00:05:30:00:06:de |
| 5 | 2.1(1a) | 4.0 | -- |

| Mod | Application Image Description | Application Image Version |
|-----|-------------------------------|---------------------------|
| 4 | SSI linecard image | 2.1(1a) |

| Mod | MAC-Address(es) | Serial-Num |
|-----|--|-------------|
| 4 | 00-05-30-00-9e-b2 to 00-05-30-00-9e-b6 | JAB06480590 |
| 5 | 00-0c-30-da-7c-18 to 00-0c-30-da-7c-1c | JAB080507P4 |

* this terminal session



Note The next time you reboot the switch, the saved image is used. If you do not save the configuration, the previously saved startup configuration image is used.

Replacing Modules

If you replace a module in a switch that contains an SSI image, you should consider the following:

- If you replace an SSM module with another SSM module, you can leave the SSI image installed on the active supervisor.
- If you replace an SSM module with any other module, you can leave the SSI image installed on the active supervisor or remove it. The supervisor module detects the module type and boots the module appropriately.

Send documentation comments to mdsfeedback-doc@cisco.com.

- If you replace a supervisor module in a switch with active and standby supervisors, no action is required because the SSI image is automatically synced to the new supervisor module.
- If you replace a supervisor module in a switch with no standby supervisor, you need to reimplement the configuration on the new supervisor.

Default Settings

Table 7-9 lists the default settings for the supervisor module.

Table 7-9 **Default Supervisor Module Settings**

| Parameters | Default |
|----------------------------|---|
| Administrative connection | Serial connection. |
| Global switch information | <ul style="list-style-type: none">• No value for system name.• No value for system contact.• No value for location. |
| System clock | No value for system clock time. |
| In-band (VSAN 1) interface | IP address, subnet mask, and broadcast address assigned to the VSAN are set to 0.0.0.0. |

Send documentation comments to mdsfeedback-doc@cisco.com.



Managing System Hardware

This chapter provides details on monitoring the health of the switch. It includes the following sections:

- [Displaying Switch Hardware Inventory, page 8-2](#)
- [Displaying the Switch Serial Number, page 8-4](#)
- [Displaying Power Usage Information, page 8-5](#)
- [Power Supply Configuration Modes, page 8-6](#)
- [About Module Temperature, page 8-9](#)
- [About Fan Modules, page 8-10](#)
- [About Clock Modules, page 8-11](#)
- [Displaying Environment Information, page 8-11](#)
- [Default Settings, page 8-12](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying Switch Hardware Inventory

Use the **show inventory** command to view information on the field replaceable units (FRUs) in the switch, including product IDs, serial numbers, and version IDs. See [Example 8-1](#).

Example 8-1 *Displays the Hardware Inventory*

```
switch# show inventory
NAME: "Chassis",  DESCR: "MDS 9506 chassis"
PID: DS-C9506           ,  VID: 0.104,  SN: FOX0712S00T

NAME: "Slot 3",  DESCR: "2x1GE IPS, 14x1/2Gbps FC Module"
PID: DS-X9302-14K9      ,  VID: 0.201,  SN: JAB081405AF

NAME: "Slot 4",  DESCR: "2x1GE IPS, 14x1/2Gbps FC Module"
PID: DS-X9302-14K9      ,  VID: 0.201,  SN: JAB081605A5

NAME: "Slot 5",  DESCR: "Supervisor/Fabric-1"
PID: DS-X9530-SF1-K9     ,  VID: 4.0,   SN: JAB0747080H

NAME: "Slot 6",  DESCR: "Supervisor/Fabric-1"
PID: DS-X9530-SF1-K9     ,  VID: 4.0,   SN: JAB0746090H

NAME: "Slot 17",  DESCR: "MDS 9506 Power Supply"
PID: DS-CAC-1900W        ,  VID: 1.0,   SN: DCA07216052

NAME: "Slot 19",  DESCR: "MDS 9506 Fan Module"
PID: DS-6SLOT-FAN        ,  VID: 0.0,   SN: FOX0638S150
```

Use the **show hardware** command to display switch hardware inventory details. See [Example 8-2](#).



Note

To display and configure modules, see [Chapter 7, “Managing Modules.”](#)

Example 8-2 *Displays the Hardware Information*

```
switch# show hardware
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2003-2004 by Cisco Systems, Inc. All rights reserved.
The copyright for certain works contained herein are owned by
Cisco Systems, Inc. and/or other third parties and are used and
distributed under license.

Software
  BIOS:      version 1.0.8
  loader:    version 1.1(0.114)
  kickstart: version 1.3(4a)
  system:    version 1.3(4a)

  BIOS compile time:      08/07/03
  kickstart image file is: bootflash:///boot-17r
  kickstart compile time: 10/25/2010 12:00:00
  system image file is:   bootflash:///isan-17r
  system compile time:    10/25/2020 12:00:00

Hardware
  RAM 1024592 kB
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
bootflash: 1000944 blocks (block size 512b)
slot0:      0 blocks (block size 512b)
```

```
172.22.90.21 uptime is 7 days 4 hours 48 minute(s) 2 second(s)
```

```
Last reset at 272247 usecs after Thu Sep 11 21:47:05 1980
Reason: Reset Requested by CLI command reload
System version: 1.3(4a)
```

```
This supervisor carries Pentium processor with 1024592 kB of memory
Intel(R) Pentium(R) III CPU at family with 512 KB L2 Cache
Rev: Family 6, Model 11 stepping 1
```

```
512K bytes of non-volatile memory.
1000944 blocks of internal bootflash (block size 512b)
```

```
-----
Chassis has 9 slots for Modules
-----
```

```
Module in slot 1 is empty
```

```
Module in slot 2 is empty
```

```
Module in slot 3 is empty
```

```
Module in slot 4 is empty
```

```
Module in slot 5 is ok
Module type is "Supervisor/Fabric-1"
No submodules are present
Model number is DS-X9530-SF1-K9
H/W version is 1.0
Part Number is 73-7523-06
Part Revision is A0
Manufacture Date is Year 6 Week 47
Serial number is JAB064705E1
CLEI code is CNP6NT0AAA
```

```
Module in slot 6 is empty
```

```
Module in slot 7 is empty
```

```
Module in slot 8 is empty
```

```
Module in slot 9 is empty
```

```
-----
Chassis has 2 Slots for Power Supplies
-----
```

```
PS in slot A is ok
Power supply type is "1153.32W 110v AC"
Model number is WS-CAC-2500W
H/W version is 1.0
Part Number is 34-1535-01
Part Revision is A0
Manufacture Date is Year 6 Week 16
Serial number is ART061600US
CLEI code is
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
PS in slot B is ok
Power supply type is "1153.32W 110v AC"
Model number is WS-CAC-2500W
H/W version is 1.0
Part Number is 34-1535-01
Part Revision is A0
Manufacture Date is Year 5 Week 41
Serial number is ART0541003V
CLEI code is
```

```
-----
Chassis has one slot for Fan Module
-----
```

```
Fan module is ok
Model number is WS-9SLOT-FAN
H/W version is 0.0
Part Number is 800-22342-01
Part Revision is
Manufacture Date is Year 0 Week 0
Serial number is
CLEI code is
```

Displaying the Switch Serial Number

The serial number of your Cisco MDS 9000 Family switch can be obtained by looking at the serial number label on the back of the switch (next to the power supply), or by executing the operating system **show sprom backplane 1** command.

```
switch# show sprom backplane 1
DISPLAY backplane sprom contents:
Common block:
Block Signature : 0xabab
Block Version   : 2
Block Length    : 156
Block Checksum  : 0x106f
EEPROM Size     : 512
Block Count     : 3
FRU Major Type  : 0x6001
FRU Minor Type  : 0x0
OEM String      : Cisco Systems, Inc.
Product Number  : DS-C9506
Serial Number   : FOX0712S007
Part Number     : 73-8697-01
Part Revision   : 01
Mfg Deviation   : 0
H/W Version     : 0.1
Mfg Bits        : 0
Engineer Use    : 0
snmpOID         : 9.12.3.1.4.26.0.0
Power Consump   : 0
RMA Code        : 0-0-0-0
Chassis specific block:
...
```



Note

If you are installing a new license, use the **show license host-id** command to obtain the switch serial. Refer to [Chapter 3, "Obtaining and Installing Licenses"](#) for further information.

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying Power Usage Information

Use the **show environment power** command to display the actual power usage information for the entire switch. In response to this command, power supply capacity and consumption information is displayed for each module. See [Example 8-3](#).



Note

In a Cisco MDS 9500 Series switch, power usage is reserved for both supervisors regardless of whether one or both supervisor modules are present.

Example 8-3 Displays Power Management Information

```
switch# show environment power
```

```
-----
PS   Model                Power      Power      Status
      (Watts)      (Amp @42V)
-----
1    DS-CAC-2500W        1153.32    27.46      ok
2    WS-CAC-2500W        1153.32    27.46      ok

Mod Model                Power      Power      Power      Power      Status
      Requested Requested Allocated Allocated
      (Watts)      (Amp @42V) (Watts)      (Amp @42V)
-----
1    DS-X9032            199.92     4.76       199.92     4.76       powered-up
4    DS-X9032            199.92     4.76       199.92     4.76       powered-up
5    DS-X9530-SF1-K9      126.00     3.00       126.00     3.00       powered-up
6    DS-X9530-SF1-K9      126.00     3.00       126.00     3.00       powered-up
9    DS-X9016            220.08     5.24       220.08     5.24       powered-up

Power Usage Summary:
-----
Power Supply redundancy mode:                redundant

Total Power Capacity                        1153.32 W

Power reserved for Supervisor(s)[-]          252.00 W
Power reserved for Fan Module(s) [-]         0.00 W
Power currently used by Modules[-]           619.92 W

Total Power Available                        281.40 W
-----
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Power Supply Configuration Modes

Switches in the MDS 9000 Family have two redundant power supply slots. The power supplies can be configured in either redundant or combined mode.

- Redundant mode—Uses the capacity of one power supply only. This is the default mode. In case of power supply failure, the entire switch has sufficient power available in the system.
- Combined mode—Uses the combined capacity of both power supplies. In case of power supply failure, the entire switch can be shut down (depends on the power used) causing traffic disruption. This mode is seldom used, except in cases where the switch has two low power supply capacities but a higher power usage.



Note

The chassis in the Cisco MDS 9000 Family uses 1200 W when powered at 110 V, and 2500 W when powered at 220 V.

To configure the power supply mode, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# power redundancy-mode combined switch(config)# | Configures combined power supply mode. |
| | switch(config)# power redundancy-mode redundant switch(config)# | Reverts to the redundant (default) power supply mode. |



Note

Use the **show environment power** command to view the current power supply configuration.

Power Supply Configuration Guidelines

Follow these guidelines when configuring power supplies:

1. When power supplies with different capacities are installed in the switch, the total power available differs based on the configured mode, either redundant or combined:

- a. Redundant mode—the total power is the lesser of the two power supply capacities. For example, suppose you have the following usage figures configured:

Power supply 1 = 2500 W

Additional power supply 2 = not used

Current usage = 2000 W

Current capacity = 2500 W

Then the following three scenarios differ as specified (see [Table 8-1](#)):

Scenario 1: If 1800 W is added as power supply 2, then power supply 2 is shut down.

Reason: 1800 W is less than the usage of 2000 W.

Scenario 2: If 2200 W is added as power supply 2, then the current capacity decreases to 2200 W.

Reason: 2200 W is the lesser of the two power supplies.

Send documentation comments to mdsfeedback-doc@cisco.com.

Scenario 3: If 3000 W is added as power supply 2, then the current capacity value remains at 2500 W.

Reason: 2500 W is the lesser of the two power supplies.

Table 8-1 Redundant Mode Power Supply Scenarios

| Scenario | Power Supply 1 (W) ¹ | Current Usage (W) | Insertion of Power Supply 2 (W) | New Capacity (W) | Action Taken by Switch |
|----------|---------------------------------|-------------------|---------------------------------|------------------|------------------------------|
| 1 | 2500 | 2000 | 1800 | 2500 | Power supply 2 is shut down. |
| 2 | 2500 | 2000 | 2200 | 2200 | Capacity becomes 2200 W. |
| 3 | 2500 | 2000 | 3300 | 2500 | Capacity remains the same. |

1. W = Watts

- b. Combined mode—the total power is twice the lesser of the two power supply capacities.

For example, suppose you have the following usage figures configured:

Power supply 1 = 2500 W

Additional Power supply 2 = not used

Current Usage = 2000 W

Current capacity = 2500 W

Then, the following three scenarios differ as specified (see [Table 8-2](#)):

Scenario 1: If 1800 W is added as power supply 2, then the capacity increases to 3600 W.

Reason: 3600 W is twice the minimum (1800 W).

Scenario 2: If 2200 W is added as power supply 2, then the current capacity increases to 4400 W.

Reason: 4400 W is twice the minimum (2200 W).

Scenario 3: If 3000 W is added as power supply 2, then the current capacity increases to 5000 W.

Reason: 5000 W is twice the minimum (2500 W).

Table 8-2 Combined Mode Power Supply Scenarios

| Scenario | Power Supply 1 (W) ¹ | Current Usage (W) | Insertion of Power Supply 2 (W) | New Capacity (W) | Action Taken by Switch |
|----------|---------------------------------|-------------------|---------------------------------|------------------|--|
| 1 | 2500 | 2000 | 1800 | 3600 | Power is never shut down. The new capacity is changed. |
| 2 | 2500 | 2000 | 2200 | 4400 | |
| 3 | 2500 | 2000 | 3300 | 5000 | |

1. W = Watts

2. When you change the configuration from combined to redundant mode and the system detects a power supply that has a capacity lower than the current usage, the power supply is shut down. If both power supplies have a lower capacity than the current system usage, the configuration is not allowed. Several configuration scenarios are summarized in [Table 8-3](#).

Send documentation comments to mdsfeedback-doc@cisco.com.

Scenario 1: You have the following usage figures configured:

Power supply 1 = 2500 W

Additional Power supply 2 = 1800 W

Current Usage = 2000 W

Current mode = combined mode (so current capacity is 3600 W)

You decide to change the switch to redundant mode. Then power supply 2 is shut down.

Reason: 1800 W is the lesser of the two power supplies and it is less than the system usage.

Scenario 2: You have the following usage figures configured:

Power supply 1 = 2500 W

Additional Power supply 2 = 2200 W

Current Usage = 2000 W

Current mode = combined mode (so current capacity is 4400 W).

You decide to change the switch to redundant mode. Then the current capacity decreases to 2200 W.

Reason: 2200 W is the lesser of the two power supplies.

Scenario 3: You have the following usage figures configured:

Power supply 1 = 2500 W

Additional Power supply 2 = 1800 W

Current Usage = 3000 W

Current mode = combined mode (so current capacity is 3600 W).

You decide to change the switch to redundant mode. Then the current capacity decreases to 2500 W and the configuration is rejected.

Reason: 2500 W is less than the system usage (3000 W).

Table 8-3 Combined Mode Power Supply Scenarios

| Scenario | Power Supply 1 (W) ¹ | Current Mode | Current Usage (W) | Power Supply 2 (W) | New Mode | New Capacity (W) | Action Taken by Switch |
|----------|---------------------------------|--------------|-------------------|--------------------|-----------|------------------|---|
| 1 | 2500 | combined | 2000 | 1800 | N/A | 3600 | This is the existing configuration. |
| | 2500 | N/A | 2000 | 1800 | redundant | 2500 | Power supply 2 is shut down |
| 2 | 2500 | combined | 2000 | 2200 | N/A | 4400 | This is the existing configuration. |
| | 2500 | N/A | 2000 | 2200 | redundant | 2200 | The new capacity is changed. |
| 3 | 2500 | combined | 3000 | 1800 | N/A | 3600 | This is the existing configuration. |
| | 2500 | N/A | 3000 | 1800 | redundant | N/A | Rejected, so the mode reverts to combined mode. |

1. W = Watts

Send documentation comments to mdsfeedback-doc@cisco.com.

About Module Temperature

Built-in, automatic sensors are provided in all switches in the Cisco MDS 9000 Family to monitor your switch at all times.

Each module (switching and supervisor) has four sensors: 1 (outlet sensor), 2 (intake sensor), 3 (onboard sensor), and 4 (onboard sensor). Each sensor has two thresholds (in degrees Celsius): minor and major.



Note

A threshold value of -127 indicates that no thresholds are configured or applicable.

- Minor threshold—When a minor threshold is exceeded, a minor alarm occurs and the following action is taken for all four sensors:
 - System messages are displayed.
 - Call Home alerts are sent (if configured).
 - SNMP notifications are sent (if configured).
- Major threshold—When a major threshold is exceeded, a major alarm occurs and the following action is taken.
 - For sensors 1, 3, and 4 (outlet and onboard sensors):

System messages are displayed.

Call Home alerts are sent (if configured).

SNMP notifications are sent (if configured).
 - For sensor 2 (intake sensor):

If the threshold is exceeded in a switching module, only that module is shut down.

If the threshold is exceeded in an active supervisor module with HA-standby or standby present, only that supervisor module is shut down and the standby supervisor module takes over.

If you do not have a standby supervisor module in your switch, you have an interval of 2 minutes to decrease the temperature. During this interval the software monitors the temperature every five (5) seconds and continuously sends system messages as configured.



Tip

To realize the benefits of these built-in, automatic sensors on any switch in the Cisco MDS 9500 Series, we highly recommend that you install dual supervisor modules. If you are using a Cisco MDS 9000 Family switch without dual supervisor modules, we recommend that you immediately replace the fan module even if one fan is not working.

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying Module Temperature

Use the **show environment temperature** command to display temperature sensors for each module (see [Example 8-4](#)).

Example 8-4 Displays Temperature Information

```
switch# show environment temperature
```

| Module | Sensor | MajorThresh (Celsius) | MinorThres (Celsius) | CurTemp (Celsius) | Status |
|--------|--------|--------------------------|-------------------------|----------------------|--------|
| 2 | Outlet | 75 | 60 | 35 | ok |
| 2 | Intake | 65 | 50 | 33 | ok |
| 5 | Outlet | 75 | 60 | 44 | ok |
| 5 | Intake | 65 | 50 | 36 | ok |
| 6 | Outlet | 75 | 60 | 42 | ok |
| 6 | Intake | 65 | 50 | 35 | ok |
| 7 | Outlet | 75 | 60 | 33 | ok |
| 7 | Intake | 65 | 50 | 30 | ok |
| 9 | Outlet | 75 | 60 | 34 | ok |
| 9 | Intake | 65 | 50 | 39 | ok |

About Fan Modules

Hot-swappable fan modules (fan trays) are provided in all switches in the Cisco MDS 9000 Family to manage airflow and cooling for the entire switch. Each fan module contains multiple fans to provide redundancy. The switch can continue functioning in the following situations:

- One or more fans fail within a fan module—Even with multiple fan failures, switches in the Cisco MDS 9000 Family can continue functioning. When a fan fails within a module, the functioning fans in the module increase their speed to compensate for the failed fan(s).
- The fan module is removed for replacement—The fan module is designed to be removed and replaced while the system is operating without presenting an electrical hazard or damage to the system. When replacing a failed fan module in a running switch, be sure to replace the new fan module within five minutes.



Tip

If one or more fans fail within a fan module, the Fan Status LED turns red. A fan failure could lead to temperature alarms if not corrected immediately.

The fan status is continuously monitored by the Cisco SAN-OS software. In case of a fan failure, the following action is taken:

- System messages are displayed.
- Call Home alerts are sent (if configured).
- SNMP notifications are sent (if configured).

Send documentation comments to mdsfeedback-doc@cisco.com.

Use the **show environment fan** command to display the fan module status (see [Example 8-5](#)).

Example 8-5 Displays Chassis Fan Information

```
switch# show environment fan
-----
FAN              Model              Hw              Status
-----
Chassis          WS-9SLOT-FAN          0.0            ok
PS-1             --                    --              ok
PS-2             --                    --              ok
```

About Clock Modules

All switches in the Cisco MDS 9000 Family have two clock modules—Module A (primary) and Module B (redundant). The clock modules are designed, tested, and qualified for mission-critical availability with a mean time between failures (MTBF) of 3,660,316 hours. This translates to a potential failure every 365 years. Additionally, Cisco MDS 9000 Family switches are designed to automatically switch to the redundant clock module should the active clock module fail.



Tip

We recommend that the failed clock module be replaced during a maintenance window.

Use the **show environment clock** command to display the status for both clock modules (see [Example 8-6](#)).

Example 8-6 Displays Chassis Clock Information

```
switch# show environment clock
-----
Clock            Model              Hw              Status
-----
A                 DS-C9500-CL          0.0            ok/active
B                 DS-C9500-CL          0.0            ok/standby
```

Displaying Environment Information

Use the **show environment** command to display all environment-related switch information.

Example 8-7 Displays All Environment Information

```
switch# show environment
Clock:
-----
Clock            Model              Hw              Status
-----
A                 Clock Module        1.0            ok/active
B                 Clock Module        1.0            ok/standby

Fan:
-----
FAN              Model              Hw              Status
-----
Chassis          DS-2SLOT-FAN        0.0            ok
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```

PS-1          --          --          ok
PS-2          --          --          absent
Temperature:
-----
Module   Sensor   MajorThresh   MinorThres   CurTemp   Status
              (Celsius)   (Celsius)   (Celsius)
-----
1         1       75           60           32        ok
1         2       65           50           32        ok
1         3      -127          -127          43        ok
1         4      -127          -127          39        ok
Power Supply:
-----
PS  Model                Power      Power      Status
              (Watts)      (Amp @42V)
-----
1   PWR-950-AC            919.38     21.89      ok
2   --                    --          --          absent
Mod Model                Power      Power      Power      Power      Status
              Requested Requested  Allocated Allocated
              (Watts)      (Amp @42V) (Watts)      (Amp @42V)
-----
1   DS-X9216-K9-SUP      220.08     5.24       220.08     5.24       powered-up
Power Usage Summary:
-----
Power Supply redundancy mode:                redundant
Total Power Capacity                        919.38 W
Power reserved for Supervisor(s) [-]         220.08 W
Power reserved for Fan Module(s) [-]          0.00 W
Power currently used by Modules [-]           0.00 W
-----
Total Power Available                        699.30 W
-----

```

Default Settings

Table 8-4 lists the default hardware settings.

Table 8-4 **Default Hardware Parameters**

| Parameters | Default |
|-------------------|-----------------|
| Power supply mode | Redundant mode. |



Using the CFS Infrastructure

As of Release 2.0(1b), the Cisco SAN-OS software uses the Cisco Fabric Services (CFS) infrastructure to enable efficient database distribution and to foster device flexibility. It simplifies SAN provisioning by automatically distributing configuration information to all switches in a fabric.

Several Cisco SAN-OS applications use the CFS infrastructure to maintain and distribute the contents of a particular application's database.

- [About CFS, page 9-2](#)
- [CFS Features, page 9-3](#)
- [Cisco SAN-OS Features Using CFS, page 9-2](#)
- [CFS Protocol, page 9-3](#)
- [CFS Distribution Scopes, page 9-4](#)
- [CFS Distribution Modes, page 9-4](#)
- [Disabling CFS Distribution on a Switch, page 9-5](#)
- [CFS Application Requirements, page 9-5](#)
- [Enabling CFS for an Application, page 9-5](#)
- [Locking the Fabric, page 9-6](#)
- [Committing Changes, page 9-6](#)
- [Discarding Changes, page 9-6](#)
- [Saving the Configuration, page 9-6](#)
- [Clearing a Locked Session, page 9-7](#)
- [CFS Merge Support, page 9-7](#)
- [Displaying CFS Configuration Information, page 9-7](#)
- [Default Settings, page 9-11](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

About CFS

Many features in the Cisco MDS switches require configuration synchronization in all switches in the fabric. Maintaining configuration synchronization across a fabric is important to maintain fabric consistency. In the absence of a common infrastructure, such synchronization is achieved through manual configuration at each switch in the fabric. This process is tedious and error prone.

As of Release 2.0(1b), Cisco Fabric Services (CFS) provides a common infrastructure for automatic configuration synchronization in the fabric. It provides the transport function as well as a rich set of common services to the applications. CFS has the ability to discover CFS capable switches in the fabric and discovering application capabilities in all CFS capable switches.

Cisco SAN-OS Features Using CFS

The following Cisco SAN-OS features use the CFS infrastructure:

- NTP (see [“NTP Configuration Distribution” section on page 4-19](#)).
- Dynamic Port VSAN Membership (see [Chapter 11, “Creating Dynamic VSANs”](#)).
- Distributed Device Alias Services (see [Chapter 16, “Distributing Device Alias Services”](#)).
- IVR topology (see [“Database Merge Guidelines” section on page 17-23](#)).
- TACACS and RADIUS (see the [“Distributing AAA Server Configuration” section on page 19-15](#)).
- User and administrator roles (see [“Role-Based Authorization” section on page 19-21](#)).
- Port security (see [“Port Security Configuration Distribution” section on page 21-9](#)).
- iSNS (see [“Configuring iSCSI Storage Name Services” section on page 28-106](#)).
- Call Home (see [“Call Home Configuration Distribution” section on page 30-12](#)).
- Syslog (see [“System Message Logging Configuration Distribution” section on page 36-7](#)).
- Fctimer (see [“fctimer Distribution” section on page 39-3](#)).
- SCSI Flow Services (see [“Enabling SCSI Flow Configuration Distribution” section on page 4](#)).
- Saving the configuration (see [“Saving the Configuration” section on page 27](#)).

Send documentation comments to mdsfeedback-doc@cisco.com.

CFS Features

CFS has the following features:

- Three modes of distribution
 - Coordinated distributions

Only one distribution is allowed in the fabric at any given time.
 - Uncoordinated distributions

Multiple parallel distributions are allowed in the fabric except when a coordinated distribution is in progress.
 - Unrestricted uncoordinated distributions

As of Cisco SAN-OS Release 2.1(1a), multiple parallel distributions are allowed in the fabric in the presence of an existing coordinated distribution. Unrestricted uncoordinated distributions are allowed to run in parallel with all other types of distributions.
- Three scopes of distribution
 - Logical scope

The distribution occurs within the scope of a VSAN.
 - Physical scope

The distribution spans the entire physical topology.
 - Over a selected set of VSANs

As of Cisco SAN-OS Release 2.1(1a), some applications, such as Inter-VSAN Routing (IVR), require configuration distribution over some specific VSANs. These applications can specify to CFS the set of VSANs over which to restrict the distribution.
- A peer-to-peer protocol that does not have a client-server relationship at the CFS layer.
- Supports a merge protocol which facilitates the merge of application configuration during a fabric merge event (when two independent fabrics merge).

CFS Protocol

The CFS functionality is independent of the lower layer transport. Currently, in Cisco MDS switches, the CFS protocol layer resides on top of the FC2 layer. CFS uses the FC2 transport services to send information to other switches. CFS uses a proprietary SW_ILS (0x77434653) protocol for all CFS packets. CFS packets are sent to/from the switch domain controller addresses.

Applications that use CFS are completely unaware of the lower layer transport.

Send documentation comments to mdsfeedback-doc@cisco.com.

CFS Distribution Scopes

Different applications on the Cisco MDS 9000 Family switches need to distribute the configuration at various levels:

- VSAN level

Applications that operate within the scope of a VSAN have the configuration distribution restricted to the VSAN. An example application is port security where the configuration database is applicable only within a VSAN.

- Physical topology level

Applications might need to distribute the configuration to the entire physical topology spanning several VSANs. Such applications include NTP and DPVM (WWN based VSAN), which are independent of VSANs.

- Between two switches

Applications might only operate between two switches in the fabric. An example application is SCSI Flow Services.

CFS Distribution Modes

CFS supports different distribution modes to support different application requirements: coordinated and uncoordinated distributions. Both modes are mutually exclusive. Only one mode is allowed at any given time.

Uncoordinated Distribution

Uncoordinated distributions are used to distribute information that is not expected to conflict with that from a peer. An example is local device registrations like in the case of iSNS. Parallel uncoordinated distributions are allowed for an application.

Coordinated Distribution

Coordinated distributions where an application can have only one such distribution at a given time. CFS uses locks to enforce this. A coordinated distribution is not allowed to start if locks are taken for the application anywhere in the fabric. A coordinated distribution consists of three stages:

1. A fabric lock is acquired.
2. The configuration is distributed and committed.
3. The fabric lock is released.

Coordinated distribution has two variants:

- CFS driven—the stages are executed by CFS in response to an application request without intervention from the application.
- Application driven—the stages are under the complete control of the application.

Coordinated distributions are used to distribute information that can be manipulated and distributed from multiple switches, for example, the port security configuration.

Send documentation comments to mdsfeedback-doc@cisco.com.

Disabling CFS Distribution on a Switch

By default, CFS distribution is enabled. Applications can distribute data and configuration information to all CFS-capable switches in the fabric where the applications exist. This is the normal mode of operation.

As of Cisco MDS SAN-OS 2.1(1a), you can globally disable CFS on a switch to isolate the applications using CFS from fabric-wide distributions while maintaining physical connectivity. When CFS is globally disabled on a switch, CFS operations are restricted to the switch and all CFS commands continue to function as if the switch were physically isolated.

To globally disable CFS distribution on a switch, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# no cfs distribution | Globally disables CFS distribution for all applications on the switch. |
| | switch(config)# cfs distribution | Enables (default) CFS distribution on the switch. |

CFS Application Requirements

All switches in the fabric must be CFS capable. A Cisco MDS switch is CFS capable, if it is running Cisco SAN-OS Release 2.0(1b) or later. Non-CFS capable switches do not receive the distribution and may result in part of the fabric not receiving the intended distribution.

CFS has the following requirements:

- **Implicit CFS Usage**—The first time you issue a CFS task for a CFS-enabled application, the configuration modification process begins and the application locks the fabric.
- **Pending database**—The pending database is a temporary buffer to hold uncommitted information. The uncommitted changes are not applied immediately to ensure that the database is synchronized with the database in the other switches in the fabric. When you commit the changes, the pending database overwrites the configuration database (also known as active database or the effective database).
- **CFS distribution is enabled or disabled on a per-application basis.** The default (enable or disable) for each application differs between applications. If CFS distribution is disabled for an application, then that application does not distribute any configuration nor does it accept a distribution from other switches in the fabric.
- **Explicit commit**—Each application requires an explicit **commit** command to commit the changes in the temporary buffer. The changes in the temporary buffer are not applied if you do not issue the **commit** command. The **commit** command distributes the new database in the fabric and then releases the fabric lock.

Enabling CFS for an Application

All CFS based applications provide an option to enable or disable the distribution capabilities. Features that existed prior to Cisco SAN-OS Release 2.0(1b) have the distribution capability disabled by default and must have their distribution capabilities enabled explicitly.

Send documentation comments to mdsfeedback-doc@cisco.com.

Applications that are introduced in Cisco SAN-OS Release 2.0(1b) have the distribution enabled by default.

The application configuration is not distributed by CFS unless distribution is explicitly enabled for that application.

Locking the Fabric

When you configure (first time configuration) a Cisco SAN-OS feature (or application) which uses the CFS infrastructure, that feature starts a CFS session and locks the fabric. When a fabric is locked, the Cisco SAN-OS software does not allow any configuration changes from a switch, other than the switch holding the lock, to this Cisco SAN-OS feature and issues a message to inform the user about the locked status. The configuration changes are held in a pending database by that application.

If you start a CFS session that requires a fabric lock but forget to end the session, an administrator can clear the session. If you lock a fabric at any time, your user name is remembered across restarts and switchovers. If another user (on the same machine) tries to perform configuration tasks, that user's attempts are rejected.

Committing Changes

A commit task saves the pending database for all application peers and releases the lock for all switches.

In general, the commit function does not start a session—only a lock function starts a session. However, an empty commit is allowed, if configuration changes are not previously made. In this case, a commit results in a session whereby locks are acquired and the current database distributed.

When you commit configuration changes to a feature using the CFS infrastructure, you receive a notification about one of the following responses:

- One or more external switches report a successful status—The application applies the changes locally and releases the fabric lock.
- None of the external switches report a successful state—The application considers this state as a failure and does not apply the changes to any switch in the fabric. The fabric lock is not released.

You can commit changes for a specified feature by using the **commit** command for that feature.

Discarding Changes

If you discard configuration changes, the application flushes the pending database and releases locks in the fabric. Both the abort and commit functions are only supported from the switch from which the fabric lock is acquired.

You can discard changes for a specified feature by using the **abort** command for that feature.

Saving the Configuration

Configuration changes that have not been applied yet (still in the pending database) are not shown in the running configuration. The configuration changes in the pending database overwrites the configuration in the effective database when you commit the changes.

Send documentation comments to mdsfeedback-doc@cisco.com.

**Caution**

If you do not commit the changes, they are not saved to the running configuration.

The CISCO-CFS-MIB contains SNMP configuration information for any CFS-related functions. Refer to the *Cisco MDS 9000 Family MIB Quick Reference* for more information on this MIB.

Clearing a Locked Session

You can clear locks held by an application from any switch in the fabric. This option is provided to rescue you from situations where locks are acquired and not released. This function requires ADMIN permissions.

**Caution**

Exercise caution when using this function to clear locks in the fabric. Any pending configurations in any switch in the fabric is flushed and lost.

CFS Merge Support

An application keeps the configuration synchronized in a fabric through CFS. Two such fabrics might merge as a result of an ISL coming up between them. These two fabrics could have two different sets of configuration information that needs to be reconciled in the event of a merge. CFS provides notification each time an application peer comes online. If two fabrics with M and N application peers merge and if a application triggers a merge action on every such notification, a link up event results in M*N merges in the fabric.

CFS supports a protocol that reduces the number of merges required to one by handling the complexity of the merge at the CFS layer. This protocol runs per application per scope. The protocol involves selecting one switch in a fabric as the merge manager for that fabric. The other switches do not play any role in the merge process.

During a merge, the merge manager in the two fabrics exchange their configuration databases with each other. The application on one of them merges the information, decides if the merge is successful and informs all switches in the combined fabric of the status of the merge.

In case of a successful merge, the merged database is distributed to all switches in the combined fabric and the entire new fabric remains in a consistent state. You can recover from a merge failure by starting a distribution from any of the switches in the new fabric. This distribution restores all peers in the fabric to the same configuration database.

Displaying CFS Configuration Information

The **show cfs status** command displays the status of CFS distribution on the switch (see [Example 9-1](#)).

Example 9-1 Displays CFS Distribution Status

```
switch# show cfs status
Fabric distribution Enabled
```

Send documentation comments to mdsfeedback-doc@cisco.com.

The **show cfs application** command displays the applications which are currently registered with CFS (see [Example 9-2](#)). The first column displays the application name. The second column indicates whether the application has registered with CFS (*enabled* or *disabled*). The last column indicates the scope of distribution for the application (*logical*, *physical*, or *both*).

Example 9-2 *Displays the Currently-Registered Applications Using CFS*

```
switch# show cfs application
-----
Application      Enabled      Scope
-----
ntp              Yes         Physical
role             No          Physical
vsan             No          Physical
radius           No          Physical
syslogd          No          Physical
callhome         No          Physical
device-alias     Yes         Physical
port-security    Yes         Logical

Total number of entries = 8
```

The **show cfs application name** command displays the details for a particular application. It displays the enabled/disabled state, timeout as registered with CFS, merge capability (if it has registered with CFS for merge support), and lastly the distribution scope. (see [Example 9-3](#))

Example 9-3 *Displays a Specified CFS Application*

```
switch# show cfs application name ntp

Enabled          : Yes
Timeout          : 5s
Merge Capable    : Yes
Scope            : Physical
```

The **show cfs lock** command displays all the locks that are currently acquired by any application. For each application the command displays the application name and scope of the lock taken. If the application lock is taken in the physical scope it indicates the switch WWN, IP address, user name, and

Send documentation comments to mdsfeedback-doc@cisco.com.

user type of the lock holder. If the application is taken in the logical scope then the VSAN in which the lock is taken, the domain, IP address, user name, and user type of the lock holder are displayed.(see [Example 9-4](#))

Example 9-4 Displays the Currently Locked Applications

```
switch# show cfs lock

Application: ntp
Scope      : Physical
-----
Switch WWN          IP Address      User Name    User Type
-----
20:00:00:05:30:00:6b:9e  10.76.100.167  admin       CLI/SNMP v3
Total number of entries = 1

Application: port-security
Scope      : Logical
-----
VSAN   Domain   IP Address      User Name    User Type
-----
1      238      10.76.100.167  admin       CLI/SNMP v3
2      211      10.76.100.167  admin       CLI/SNMP v3
Total number of entries = 2
```

The **show cfs lock name** command displays the lock details similar for the specified application.(see [Example 9-5](#))

Example 9-5 Displays the Lock Information for the Specified Application

```
switch# show cfs lock name ntp
Scope      : Physical
-----
Switch WWN          IP Address      User Name    User Type
-----
20:00:00:05:30:00:6b:9e  10.76.100.167  admin       CLI/SNMP v3

Total number of entries = 1
```

The **show cfs merge status name** command displays the merge status for a given application. [Example 9-6](#) displays the output for an application distributing in logical scope shows the merge status in all the valid VSANs on the switch. The command shows the scope of merge i.e. logical and the vsan. It shows the merge status which is one of the following Success, waiting, Failure Or In Progress. In case of a successful merge all the switches in the fabric are shown under the local fabric. In case of a merge failure or a merge being in progress the local fabric and the remote fabric involved in the merge are indicated separately. The application server in each fabric which is mainly responsible for the merge is indicated by the term Merge Master.

Example 9-6 Displays the Merge Status for the Specified Application

```
switch# show cfs merge status name port-security

Logical [VSAN 1] Merge Status: Failed
Local Fabric
-----
Domain Switch WWN          IP Address
-----
238      20:00:00:05:30:00:6b:9e  10.76.100.167 [Merge Master]
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```

Remote Fabric
-----
Domain Switch WWN          IP Address
-----
236      20:00:00:0e:d7:00:3c:9e 10.76.100.169 [Merge Master]

Logical [VSAN 2] Merge Status: Success
Local Fabric
-----
Domain Switch WWN          IP Address
-----
211      20:00:00:05:30:00:6b:9e 10.76.100.167 [Merge Master]
1        20:00:00:0e:d7:00:3c:9e 10.76.100.169

Logical [VSAN 3] Merge Status: Success
Local Fabric
-----
Domain Switch WWN          IP Address
-----
221      20:00:00:05:30:00:6b:9e 10.76.100.167 [Merge Master]
103      20:00:00:0e:d7:00:3c:9e 10.76.100.169

```

The **show cfs merge status name** command displayed in [Example 9-7](#) shows an application using the physical scope with a merge failure. The command uses the specified application name to display the merge status based on the application scope.

Example 9-7 *Displays the Merge Status for a Physical Scope with Merge Failure*

```

switch# show cfs merge status name ntp

Physical Merge Status: Failed
Local Fabric
-----
Switch WWN          IP Address
-----
20:00:00:05:30:00:6b:9e 10.76.100.167 [Merge Master]

Remote Fabric
-----
Switch WWN          IP Address
-----
20:00:00:0e:d7:00:3c:9e 10.76.100.169 [Merge Master]

```

The **show cfs peers** command displays all the switches in the physical fabric in terms of the switch WWN and the IP address. The local switch is indicated as `Local` (see [Example 9-8](#)).

Example 9-8 *Displays Peers in the Fabric*

```

switch# show cfs peers

Physical Fabric
-----
Switch WWN          IP Address
-----
20:00:00:05:30:00:6b:9e 10.76.100.167 [Local]
20:00:00:0e:d7:00:3c:9e 10.76.100.169

Total number of entries = 2

```

Send documentation comments to mdsfeedback-doc@cisco.com.

The **show cfs peers name** command displays all the peers for which a particular application is registered with CFS. The output shows all the peers for physical scope or for each of the valid vsans on the switch depending on the application scope. For physical scope the switch WWN for all the peers are indicated. The local switch is indicated as `Local` (see [Example 9-9](#)).

Example 9-9 Displays Peers for a Specified CFS-Registered Application

```
switch# show cfs peers name ntp

Scope      : Physical
-----
Switch WWN      IP Address
-----
20:00:00:44:22:00:4a:9e  172.22.92.27    [Local]
20:00:00:05:30:01:1b:c2  172.22.92.215
```

The **show cfs peers name** command in [Example 9-10](#) displays all the application peers (all switches in which that application is registered). The local switch is indicated as `Local`.

Example 9-10 Displays Logical Scope for Each VSAN

```
switch# show cfs peers name port-security
Scope      : Logical [VSAN 1]
-----
Domain      Switch WWN      IP Address
-----
124         20:00:00:44:22:00:4a:9e  172.22.92.27    [Local]
98          20:00:00:05:30:01:1b:c2  172.22.92.215

Total number of entries = 2

Scope      : Logical [VSAN 3]
-----
Domain      Switch WWN      IP Address
-----
224         20:00:00:44:22:00:4a:9e  172.22.92.27    [Local]
151         20:00:00:05:30:01:1b:c2  172.22.92.215

Total number of entries = 2
```

Default Settings

[Table 9-1](#) lists the default settings for CFS configurations.

Table 9-1 Default CFS Parameters

| Parameters | Default |
|--------------------------|---|
| Database changes | Implicitly enabled with the first configuration change. |
| Application distribution | Differs based on application. |
| Commit | Explicit configuration is required. |

Send documentation comments to mdsfeedback-doc@cisco.com.



Configuring and Managing VSANs

You can achieve higher security and greater stability in Fibre Channel fabrics by using virtual SANs (VSANs). VSANs provide isolation among devices that are physically connected to the same fabric. With VSANs you can create multiple logical SANs over a common physical infrastructure. Each VSAN can contain up to 239 switches and has an independent address space that allows identical Fibre Channel IDs (FCIDs) to be used simultaneously in different VSANs. This chapter includes the following sections:

- [VSAN Advantages, page 10-1](#)
- [How VSANs Work, page 10-2](#)
- [VSANs Versus Zones, page 10-4](#)
- [Default and Isolated VSANs, page 10-5](#)
- [VSAN Attributes, page 10-6](#)
- [VSAN Membership, page 10-6](#)
- [Creating and Configuring VSANs Statically, page 10-7](#)
- [Default Settings, page 10-10](#)

VSAN Advantages

VSANs offer the following advantages:

- **Traffic isolation**—Traffic is contained within VSAN boundaries and devices reside only in one VSAN ensuring absolute separation between user groups, if desired.
- **Scalability**—VSANs are overlaid on top of a single physical fabric. The ability to create several logical VSAN layers increases the scalability of the SAN.
- **Per VSAN fabric services**—Replication of fabric services on a per VSAN basis provides increased scalability and availability.
- **Redundancy**—Several VSANs created on the same physical SAN ensure redundancy. If one VSAN fails, redundant protection (to another VSAN in the same physical SAN) is configured using a backup path between the host and the device.
- **Ease of configuration**—Users can be added, moved, or changed between VSANs without changing the physical structure of a SAN. Moving a device from one VSAN to another only requires configuration at the port level, not at a physical level.

Send documentation comments to mdsfeedback-doc@cisco.com.

How VSANs Work

A VSAN is a virtual storage area network (SAN). A SAN is a dedicated network that interconnects hosts and storage devices primarily to exchange SCSI traffic. In SANs you use the physical links to make these interconnections. A set of protocols run over the SAN to handle routing, naming, and zoning. You can design multiple SANs with different topologies.

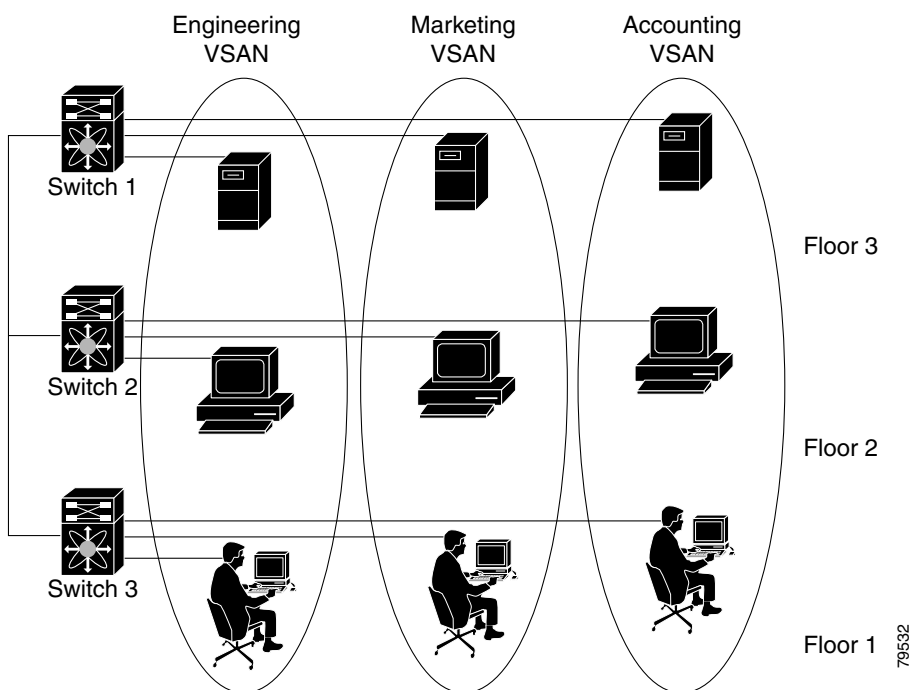
With the introduction of VSANs, the network administrator can build a single topology containing switches, links, and one or more VSANs. Each VSAN in this topology has the same behavior and property of a SAN. A VSAN has the following additional features:

- Multiple VSANs can share the same physical topology.
- The same Fibre Channel IDs (FCIDs) can be assigned to a host in another VSAN, thus increasing VSAN scalability.
- Every instance of a VSAN runs all required protocols such as FSPF, domain manager, and zoning.
- Fabric-related configurations in one VSAN do not affect the associated traffic in another VSAN.
- Events causing traffic disruptions in one VSAN are contained within that VSAN and are not propagated to other VSANs.

As displayed in both [Figure 10-1](#) and [Figure 10-2](#), the switch icons indicate that these features apply to any switch in the Cisco MDS 9000 Family.

[Figure 10-1](#) shows a fabric with three switches, one on each floor. The geographic location of the switches and the attached devices is independent of their segmentation into logical VSANs. Between VSANs no communication is possible. Within each VSAN, all members can talk to one another.


Figure 10-1 Logical VSAN Segmentation



Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 10-2 shows a physical Fibre Channel switching infrastructure with two defined VSANs: VSAN 2 (dashed) and VSAN 7 (solid). VSAN 2 includes hosts H1 and H2, application servers AS2 and AS3, and storage arrays SA1 and SA4. VSAN 7 connects H3, AS1, SA2, and SA3.

Figure 10-2 Example of two VSANs



The four switches in this network are interconnected by trunk links that carry both VSAN 2 and VSAN 7 traffic. Thus the inter-switch topology of both VSAN 2 and VSAN 7 are identical. This is not a requirement and a network administrator can enable certain VSANs on certain links to create different VSAN topologies.

Without VSANs, a network administrator would need separate switches and links for separate SANs. By enabling VSANs, the same switches and links may be shared by multiple VSANs. VSANs allow SANs to be built on port granularity instead of switch granularity. Figure 10-2 illustrates that a VSAN is a group of hosts or storage devices that communicate with each other using a virtual topology defined on the physical SAN.

The criteria for creating such groups differ based on the VSAN topology:

- VSANs can separate traffic based on the following requirements:
 - Different customers in storage provider data centers
 - Production or test in an enterprise network
 - Low and high security requirements
 - Backup traffic on separate VSANs
 - Replicating data from user traffic
- VSANs can meet the needs of a particular department or application.

Send documentation comments to mdsfeedback-doc@cisco.com.

VSANs Versus Zones

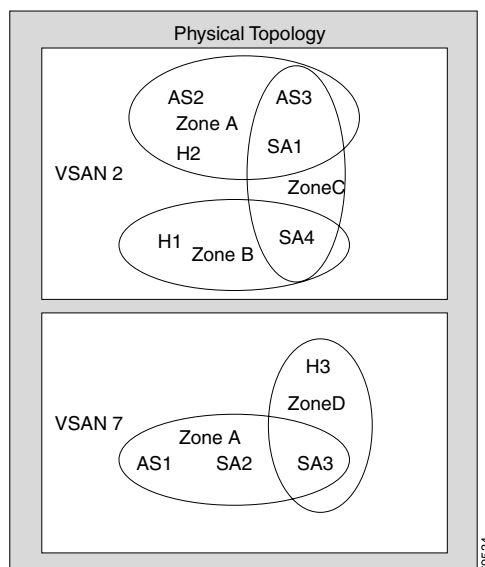
You can define multiple zones in a VSAN. Because two VSANs are equivalent to two unconnected SANs, zone A on VSAN 1 is different and separate from zone A in VSAN 2. Table 10-1 lists the differences between VSANs and zones.

Table 10-1 VSAN and Zone Comparison

| VSAN Characteristic | Zone Characteristic |
|---|---|
| VSANs equal SANs with routing, naming, and zoning protocols. | Routing, naming, and zoning protocols are not available on a per-zone basis. |
| — | Zones are always contained within a VSAN. Zones never span two VSANs. |
| VSANs limit unicast, multicast, and broadcast traffic. | Zones limit unicast traffic. |
| Membership is typically defined using the VSAN ID to Fx ports. | Membership is typically defined by the pWWN. |
| An HBA or a storage device can belong only to a single VSAN—the VSAN associated with the Fx port. | An HBA or storage device can belong to multiple zones. |
| VSANs enforce membership at each E port, source port, and destination port. | Zones enforce membership only at the source and destination ports. |
| VSANs are defined for larger environments (storage service providers). | Zones are defined for a set of initiators and targets not visible outside the zone. |
| VSANs encompass the entire fabric. | Zones are configured at the fabric edge. |

Figure 10-3 shows the possible relationships between VSANs and zones. In VSAN 2, three zones are defined: zone A, zone B, and zone C. Zone C overlaps both zone A and zone B as permitted by Fibre Channel standards. In VSAN 7, two zones are defined: zone A and zone D. No zone crosses the VSAN boundary—they are completely contained within the VSAN. Zone A defined in VSAN 2 is different and separate from zone A defined in VSAN 7.

Figure 10-3 VSANS with Zoning



Send documentation comments to mdsfeedback-doc@cisco.com.

Default and Isolated VSANs

Up to 256 VSANs can be configured in a switch. Of these, one is a default VSAN (VSAN 1), and another is an isolated VSAN (VSAN 4094). User-specified VSAN IDs range from 2 to 4093.

Default VSAN

The factory settings for switches in the Cisco MDS 9000 Family have only the default VSAN 1 enabled. If you do not need more than one VSAN for a switch, use this default VSAN as the implicit parameter during configuration. If no VSANs are configured, all devices in the fabric are considered part of the default VSAN. By default, all ports are assigned to the default VSAN.

**Note**

VSAN 1 cannot be deleted, but it can be suspended.

Isolated VSAN

VSAN 4094 is an isolated VSAN. All non-trunking ports are transferred to this VSAN when the VSAN to which they belong is deleted. This avoids an implicit transfer of ports to the default VSAN or to another configured VSAN. All ports in the deleted VSAN are isolated (disabled).

**Note**

When you configure a port in VSAN 4094 or move a port to VSAN 4094, that port is immediately isolated.

**Caution**

Do not use an isolated VSAN to configure ports.

Displaying Isolated VSAN Membership

The **show vsan 4094 membership** command displays all ports associated with the isolated VSAN.

Send documentation comments to mdsfeedback-doc@cisco.com.

VSAN Attributes

VSANs have the following attributes:

- **VSAN ID**—The VSAN ID identifies the VSAN as the default VSAN (VSAN 1), user-defined VSANs (VSAN 2 to 4093), and the isolated VSAN (VSAN 4094).
- **State**—The administrative state of a VSAN can be configured to an active (default) or suspended state. Once VSANs are created, they may exist in various conditions or states.
 - The active state of a VSAN indicates that the VSAN is configured and enabled. By enabling a VSAN, you activate the services for that VSAN.
 - The suspended state of a VSAN indicates that the VSAN is configured but not enabled. If a port is configured in this VSAN, it is disabled. Use this state to deactivate a VSAN without losing the VSAN's configuration. All ports in a suspended VSAN are disabled. By suspending a VSAN, you can preconfigure all the VSAN parameters for the whole fabric and activate the VSAN immediately.
- **VSAN name**—This text string identifies the VSAN for management purposes. The name can be from 1 to 32 characters long and it must be unique across all VSANs. By default, the VSAN name is a concatenation of VSAN and a four-digit string representing the VSAN ID. For example, the default name for VSAN 3 is VSAN0003.



Note

A VSAN name must be unique.

- **Load balancing attributes**—These attributes indicate the use of the source-destination ID (src-dst-id) or the originator exchange OX ID (src-dst-ox-id, the default) for load balancing path selection.

Operational State of a VSAN

A VSAN is in the operational state if the VSAN is active and at least one port is up. This state indicates that traffic can pass through this VSAN. This state cannot be configured.

VSAN Membership

Port VSAN membership on the switch is assigned on a port-by-port basis. By default each port belongs to the default VSAN. You can assign VSAN membership to ports using one of two methods:

- **Statically**—by assigning VSANs to ports. This method was available before Release 2.0(1b) and continues to be available in Release 2.0(1b) and later.

You can change the VSAN membership by using the **vsan number interface type slot/port** command. See the [“Creating and Configuring VSANs Statically”](#) section on page 10-7.

- **Dynamically**—by assigning VSANs based on the device WWN. This method is referred to as the Dynamic Port VSAN Membership (DPVM) feature.

See [Chapter 11, “Creating Dynamic VSANs.”](#)

Trunking ports have an associated list of VSANs that are part of an allowed list (see [Chapter 13, “Configuring Trunking”](#)).

Send documentation comments to mdsfeedback-doc@cisco.com.

Creating and Configuring VSANs Statically

You cannot configure any application-specific parameters for a VSAN before creating the VSAN.

To create and configure VSANs, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# vsan database switch(config-vsan-db)# | Configures the database for a VSAN. Application specific VSAN parameters cannot be configured from this prompt. |
| Step 3 | switch(config-vsan-db)# vsan 2 switch(config-vsan-db)# | Creates a VSAN with the specified ID (2) if that VSAN does not exist already. |
| | switch(config-vsan-db)# vsan 2 name TechDoc updated vsan 2 switch(config-vsan-db)# | Updates the VSAN with the assigned name (TechDoc). |
| Step 4 | switch(config-vsan-db)# vsan 2 loadbalancing src-dst-id switch(config-vsan-db)# | Enables the load balancing guarantee for the selected VSAN and directs the switch to use the source and destination ID for its path selection process. |
| | switch(config-vsan-db)# no vsan 2 loadbalancing src-dst-id switch(config-vsan-db)# | Negates the command issued in the previous step and reverts to the default values of the load-balancing parameters. |
| | switch(config-vsan-db)# vsan 2 loadbalancing src-dst-ox-id switch(config-vsan-db)# | Changes the path selection setting to use the source ID, the destination ID, and the OX ID (default). |
| Step 5 | switch(config-vsan-db)# vsan 2 suspend switch(config-vsan-db)# | Suspends the selected VSAN. |
| | switch(config-vsan-db)# no vsan 2 suspend vs.-config-vsan-db# | Negates the suspend command issued in the previous step. |
| Step 6 | switch(config-vsan-db)# end switch# | Returns you to EXEC mode. |

Assigning Static Port VSAN Membership

To statically assign VSAN membership for an interface, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# vsan database switch(config-vsan-db)# | Configures the database for a VSAN. |
| Step 3 | switch(config-vsan-db)# vsan 2 switch(config-vsan-db)# | Creates a VSAN with the specified ID (2) if that VSAN does not exist already. |
| Step 4 | switch(config-vsan-db)# vsan 2 interface fc1/8 switch(config-vsan-db)# | Assigns the membership of the fc1/8 interface to the specified VSAN (VSAN 2). |
| Step 5 | switch(config-vsan-db)# vsan 7 switch(config-vsan-db)# | Creates another VSAN with the specified ID (7) if that VSAN does not exist already. |

Send documentation comments to mdsfeedback-doc@cisco.com.

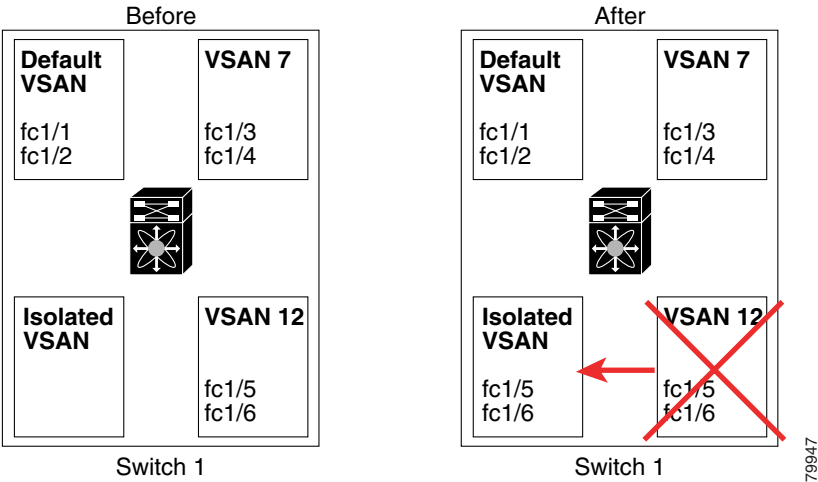
| | Command | Purpose |
|--------|--|--|
| Step 6 | <pre>switch(config-vsan-db) # vsan 7 interface fc1/8 switch(config-vsan-db) #</pre> | Updates the membership information of the interface to reflect the changed VSAN. |
| Step 7 | <pre>switch(config-vsan-db) # no vsan 7 interface fc1/8 switch(config-vsan-db) #</pre> | Removes the interface from the VSAN. |

Deleting Static VSANs

When an active VSAN is deleted, all of its attributes are removed from the running configuration. VSAN-related information is maintained by the system software as follows:

- VSAN attributes and port membership details are maintained by the VSAN manager. This feature is affected when you delete a VSAN from the configuration. When a VSAN is deleted all the ports in that VSAN are made inactive and the ports are moved to the isolated VSAN. If the same VSAN is recreated, the ports do not automatically get assigned to that VSAN. You must explicitly reconfigure the port VSAN membership (see [Figure 10-4](#)).

Figure 10-4 VSAN Port Membership Details



- VSAN-based runtime (name server), zoning, and configuration (static routes) information is removed when the VSAN is deleted.
- Configured VSAN interface information is removed when the VSAN is deleted.


Note

The allowed VSAN list is not affected when a VSAN is deleted (see [Chapter 13, “Configuring Trunking”](#)).

Any commands for a nonconfigured VSAN are rejected. For example, if VSAN 10 is not configured in the system, then a command request to move a port to VSAN 10 is rejected.

Send documentation comments to mdsfeedback-doc@cisco.com.

To delete a VSAN and its various attributes, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# vsan database switch(config-db)# | Configures the VSAN database. |
| Step 3 | switch-config-db# vsan 2 switch(config-vsan-db)# | Places you in VSAN configuration mode. |
| Step 4 | switch(config-vsan-db)# no vsan 5 switch(config-vsan-db)# | Deletes VSAN 5 from the database and switch. |
| Step 5 | switch(config-vsan-db)# end switch# | Places you in EXEC mode. |

Displaying Static VSAN Configurations

Use the **show vsan** command to display information about configured VSANs (see Examples 10-1 to 10-6).

Example 10-1 *Displays the Configuration for a Specific VSAN*

```
switch# show vsan 100
vsan 100 information
      name:VSAN0100 state:active
      in-order guarantee:no interoperability mode:no
      loadbalancing:src-id/dst-id/oxid
```

Example 10-2 *Displays the VSAN Usage*

```
switch# show vsan usage
4 vsan configured
configured vsans:1-4
vsans available for configuration:5-4093
```

Example 10-3 *Displays All VSANs*

```
switch# show vsan
vsan 1 information
      name:VSAN0001 state:active
      in-order guarantee:no interoperability mode:no
      loadbalancing:src-id/dst-id/oxid
vsan 2 information
      name:VSAN0002 state:active
      in-order guarantee:no interoperability mode:no
      loadbalancing:src-id/dst-id/oxid
vsan 7 information
      name:VSAN0007 state:active
      in-order guarantee:no interoperability mode:no
      loadbalancing:src-id/dst-id/oxid
vsan 100 information
      name:VSAN0100 state:active
      in-order guarantee:no interoperability mode:no
      loadbalancing:src-id/dst-id/oxid
vsan 4094:isolated vsan
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 10-4 Displays Membership Information for the Specified VSAN

```
switch # show vsan 1 membership
vsan 1 interfaces:
    fc1/1    fc1/2    fc1/3    fc1/4    fc1/5    fc1/6    fc1/7    fc1/9
    fc1/10   fc1/11   fc1/12   fc1/13   fc1/14   fc1/15   fc1/16   port-channel 99
```



Note

Interface information is not displayed if interfaces are not configured on this VSAN.

Example 10-5 Displays Static Membership Information for All VSANs

```
switch # show vsan membership
vsan 1 interfaces:
    fc2/16   fc2/15   fc2/14   fc2/13   fc2/12   fc2/11   fc2/10   fc2/9
    fc2/8    fc2/7    fc2/6    fc2/5    fc2/4    fc2/3    fc2/2    fc2/1
    fc1/16   fc1/15   fc1/14   fc1/13   fc1/12   fc1/11   fc1/10   fc1/9
    fc1/7    fc1/6    fc1/5    fc1/4    fc1/3    fc1/2    fc1/1
vsan 2 interfaces:
vsan 7 interfaces:
    fc1/8
vsan 100 interfaces:
vsan 4094(isolated vsan) interfaces:
```

Example 10-6 Displays Static Membership Information for a Specified Interface

```
switch # show vsan membership interface fc1/1
fc1/1
    vsan:1
    allowed list:1-4093
```

Default Settings

Table 10-2 lists the default settings for all configured VSANs.

Table 10-2 Default VSAN Parameters

| Parameters | Default |
|--------------------------|--|
| Default VSAN | VSAN 1. |
| State | Active state. |
| Name | Concatenation of VSAN and a four-digit string representing the VSAN ID. For example, VSAN 3 is VSAN0003. |
| Load-balancing attribute | OX ID (src-dst-ox-id). |



CHAPTER 11

Creating Dynamic VSANs

Port VSAN membership on the switch is assigned on a port-by-port basis. By default each port belongs to the default VSAN.

As of Cisco SAN-OS Release 2.0(1b), you can dynamically assign VSAN membership to ports by assigning VSANs based on the device WWN. This method is referred to as the Dynamic Port VSAN Membership (DPVM) feature. DPVM offers flexibility and eliminates the need to reconfigure the port VSAN membership to maintain fabric topology when a host or storage device connection is moved between two Cisco MDS switches or two ports within a switch. It retains the configured VSAN regardless of where a device is connected or moved. To assign VSANs statically, see [Chapter 10, “Configuring and Managing VSANs.”](#)

This chapter includes the following sections:

- [About DPVM, page 11-2](#)
- [DPVM Requirements, page 11-2](#)
- [Enabling DPVM, page 11-2](#)
- [About DPVM Databases, page 11-3](#)
- [Configuring Config and Pending Databases, page 11-3](#)
- [Activating Config Databases, page 11-3](#)
- [About Autolearned Entries, page 11-4](#)
- [Enabling Autolearning, page 11-4](#)
- [Configuring DPVM Database Distribution, page 11-5](#)
- [Database Merge Guidelines, page 11-7](#)
- [Copying DPVM Databases, page 11-7](#)
- [Comparing Database Differences, page 11-8](#)
- [Displaying DPVM Configurations, page 11-8](#)
- [Sample DPVM Configuration, page 11-10](#)
- [Default Settings, page 11-12](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

About DPVM

DPVM configurations are based on port world wide name (pWWN) and node world wide name (nWWN) assignments. A DPVM database contains mapping information for each device pWWN/nWWN assignment and the corresponding VSAN. The Cisco SAN-OS software checks the database during a device FLOGI and obtains the required VSAN details.

The pWWN identifies the host or device and the nWWN identifies a node consisting of multiple devices. You can assign any one of these identifiers or any combination of these identifiers to configure DPVM mapping. If you assign a combination, then preference is given to the pWWN.

DPVM uses the Cisco Fabric Services (CFS) infrastructure to enable efficient database management and distribution. DPVM uses the application driven, coordinated distribution mode and the fabric-wide distribution scope ([Chapter 9, “Using the CFS Infrastructure”](#)).

DPVM Requirements

- To use the DPVM feature as designed, be sure to verify the following requirements:
- The interface through which the dynamic device connects to the Cisco MDS 9000 Family switch must be configured as an F port.
 - The static port VSAN of the F port should be valid (not isolated, not suspended and in existence).
 - The dynamic VSAN configured for the device in the DPVM database should be valid (not isolated, not suspended and in existence).


Note

The DPVM feature overrides any existing static port VSAN membership configuration. If the VSAN corresponding to the dynamic port is deleted or suspended, the port is shut down.

Enabling DPVM

To begin configuring the DPVM feature, you must explicitly enable DPVM on the required switches in the fabric. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family.

The configuration and verification commands for the DPVM feature are only available when DPVM is enabled on a switch. When you disable this feature, all related configurations are automatically discarded.

To enable DPVM on any participating switch, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# dpvm enable | Enables DPVM on that switch. |
| | switch(config)# no dpvm enable | Disables (default) DPVM on that switch. |

Send documentation comments to mdsfeedback-doc@cisco.com.

About DPVM Databases

The DPVM database consists of a series of device mapping entries. Each entry consists of a device pWWN/nWWN assignment along with the dynamic VSAN to be assigned. You can configure a maximum of 16,000 DPVM entries in the DPVM database. This database is global to the whole switch (and fabric) and is not maintained for each VSAN.

The DPVM feature uses three databases to accept and implement configurations.

- Configuration (config) database—All configuration changes are stored in the configuration database when distribution is disabled.
- Active database—The database currently enforced by the fabric.
- Pending database—All configuration changes are stored in the pending database when distribution is enabled (see the [“Configuring DPVM Database Distribution”](#) section on page 11-5).

Changes to the config database are not reflected in the active database until you activate the config database. Changes to the pending database are not reflected in the config/active database until you commit the pending database. This database structure allows you to create multiple entries, review changes, and let the config and pending databases take effect.

Configuring Config and Pending Databases

To create and populate the config and pending databases, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# dpvm database switch(config-dpvm-db)# | Creates the config database. |
| | switch(config)# no dpvm database | Deletes the config database. |
| Step 3 | switch(config-dpvm-db)# pwwn 12:33:56:78:90:12:34:56 vsan 100 | Maps the specified device pWWN to VSAN 100. |
| | switch(config-dpvm-db)# no pwwn 12:33:56:78:90:12:34:56 vsan 101 | Removes the specified device pWWN mapping from the config database. |
| Step 4 | switch(config-dpvm-db)# nwwn 14:21:30:12:63:39:72:81 vsan 101 | Maps the specified device nWWN to VSAN 101. |
| | switch(config-dpvm-db)# no nwwn 14:21:30:12:63:39:72:80 vsan 101 | Removes the specified device nWWN mapping from the config database. |

Activating Config Databases

When you explicitly activate the config database, the config database becomes the active database. Activation may fail if conflicting entries are found between the config database and the currently active database. However, you can force activation to override conflicting entries.

To disable DPVM, you must explicitly deactivate the currently active DPVM database by issuing the **no dpvm activate** command.

Send documentation comments to mdsfeedback-doc@cisco.com.

To activate the DPVM config database, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# dpvm activate | Activates the config database. |
| | switch(config)# no dpvm activate | Deactivates the currently active database. |
| | switch(config)# dpvm activate force | Forcefully activates the config database to override conflicting entries. |

About Autolearned Entries

The DPVM database can be configured to automatically learn (autolearn) about new devices within each VSAN. The autolearn feature can be enabled or disabled at any time. Learned entries are created by populating device pWWNs and VSANs in the DPVM active database. The DPVM active database should already be available to enable the autolearn feature.

You can delete any learned entry from the DPVM active database when you enable the autolearn feature. These entries only become permanent in the active database when you disable the autolearn feature.

The following conditions apply to learned entries:

- If a device logs out while autolearn is enabled, that entry is automatically deleted from the DPVM active database.
- If the same device logs multiple times into the switch through different ports, then the VSAN corresponding to last login is remembered.
- Learned entries do not override previously configured and activated entries.
- Learning is a two-part process—enabling autolearning followed by disabling autolearning. When the **auto-learn** option is enabled, the following applies:
 - Learning currently logged-in devices—occurs from the time learning is enabled.
 - Learning new device logins— occurs as and when new devices log in to the switch.

Enabling Autolearning

To enable autolearning, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# dpvm auto-learn | Enables learning on this switch. |
| | switch(config)# no dpvm auto-learn | Disables (default) learning on this switch. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Clearing Learned Entries

You can clear DPVM entries from the active database (if autolearn is still enabled) using one of two methods.

- To clear a single autolearn entry, use the **clear dpvm auto-learn pwwn** command.

```
switch# clear dpvm auto-learn pwwn 55:22:33:44:55:66:77:88
```

- To clear all autolearn entries, use the **clear dpvm auto-learn** command.

```
switch# clear dpvm auto-learn
```



Note

These two commands do not start a session and can only be issued in the local switch.

Configuring DPVM Database Distribution

If the DPVM database is available on all switches in the fabric, devices can be moved anywhere and offer the greatest flexibility. To enable database distribution to the neighboring switches, the database should be consistently administered and distributed across all switches in the fabric. The Cisco SAN-OS software uses the Cisco Fabric Services (CFS) infrastructure to achieve this requirement (see [Chapter 9, “Using the CFS Infrastructure”](#)).

Using the CFS infrastructure, each DPVM server learns the DPVM database from each of its neighboring switches during the ISL bring-up process. If you change the database locally, the DPVM server notifies its neighboring switches, and that database is updated by all switches in the Fabric.

If fabric distribution is enabled, all changes to the configuration database are stored in the pending database. These changes include the following tasks:

- Adding, deleting, or modifying database entries.
- Activating, deactivating, or deleting the configuration database.
- Enabling or disabling autolearning.

These changes are distributed to all switches in a fabric when you commit the changes. You can also discard (abort) the changes at this point.



Tip

You can view the contents of the pending database by issuing the **show dpvm pending** command.

Disabling DPVM Database Distribution

To disable DPVM database distribution to the neighboring switches, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# no dpvm distribute | Disables DPVM distribution to the neighboring switches. |
| | switch(config)# dpvm distribute | Enables (default) DPVM distribution to the neighboring switches. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Locking the Fabric

The first action that modifies the existing configuration creates the pending database and locks the feature in the fabric. Once you lock the fabric, the following conditions apply:

- No other user can make any configuration changes to this feature.
- A copy of the configuration database becomes the pending database. Modifications from this point on are made to the pending database. The pending database remains in effect until you commit the modifications to the pending database or discard (abort) the changes to the pending database.

To lock the fabric and apply changes to the pending database, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# dpvm database switch(config-dpvm-db)# | Accesses the DPVM config database. |
| Step 3 | switch(config-dpvm-db)# pwwn 11:22:33:44:55:66:77:88 vsan 11 | Adds one entry to the DPVM config database. |
| Step 4 | switch(config-dpvm-db)# exit switch(config)# | Exits to configuration mode. |
| Step 5 | switch(config)# dpvm activate | Activates the config database. |

Committing Changes

If you commit the changes made to the configuration, the configuration in the pending database are distributed to other switches. On a successful commit, the configuration change is applied throughout the fabric and the lock is released.

To commit the pending database, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# dpvm commit | Commits the database entries that are currently in the pending database. |

Discarding Changes

If you discard (abort) the changes made to the pending database, the configurations remains unaffected and the lock is released.

To discard the pending database, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# dpvm abort | Discards the database entries that are currently in the pending database. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Clearing a Locked Session

If you have performed a DPVM task and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.



Tip

The pending database is only available in the volatile directory and are subject to being discarded if the switch is restarted.

To use administrative privileges and release a locked DPVM session, use the **clear dpvm session** command in EXEC mode.

```
switch# clear dpvm session
```

Database Merge Guidelines

A database merge refers to a union of the configuration database and static (unlearned) entries in the active database. See the “[CFS Merge Support](#)” section on page 9-7 for detailed concepts.

When merging the database between two fabric, follow these guidelines:

- Verify that the activation status and the auto-learn status is the same in both fabrics.
- Verify that the combined number of device entries in each database does not exceed 16K.



Caution

If you do not follow these two conditions, the merge will fail. The next distribution will forcefully synchronize the databases and the activation states in the fabric.

Copying DPVM Databases

The following circumstances may require the active database to be copied to the config database:

- If the learned entries are only added to the active database.
- If the config database or entries in the config database are accidentally deleted.



Note

If you copy the DPVM database and fabric distribution is enabled, you must commit the changes.

To copy the currently active database to the config database, use the **dpvm database copy** command.

```
switch# dpvm database copy active
```

Legend: "+" New Entry, "-" Missing Entry, "*" Possible Conflict Entry

```
-----
- pwnn 12:33:56:78:90:12:34:56 vsan 100
- nwnn 14:21:30:12:63:39:72:81 vsan 101
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Comparing Database Differences

You can compare the DPVM databases as follows:

- Use the **dpvm database diff active** command to compare the active database with the config database.


```
switch# dpvm database diff active
Legend: "+" New Entry, "-" Missing Entry, "*" Possible Conflict Entry
-----
- pwwn 44:22:33:44:55:66:77:88 vsan 44
* pwwn 11:22:33:44:55:66:77:88 vsan 11
```
- Use the **dpvm database diff config** command to compare config database with the active database.


```
switch# dpvm database diff config
Legend: "+" New Entry, "-" Missing Entry, "*" Possible Conflict Entry
-----
+ pwwn 44:22:33:44:55:66:77:88 vsan 44
* pwwn 11:22:33:44:55:66:77:88 vsan 22
```
- Use the **show dpvm pending-diff** command (when CFS distribution is enabled) to compare the pending database with the config database.

To add pending database entries to the DPVM config database, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# dpvm distribute | Enables CFS distribution. |
| Step 3 | switch(config)# dpvm database | Accesses the DPVM config database. |
| Step 4 | switch(config-dpvm-db) # pwwn 44:22:33:44:55:66:77:88 vsan 55 switch(config-dpvm-db) # pwwn 55:22:33:44:55:66:77:88 vsan 55 | Adds two entries to the DPVM config database. |

Displaying DPVM Configurations

Use the **show dpvm** command to display information about WWNs configured on a per VSAN basis (see Examples 11-1 to 11-6).

Example 11-1 *Displays the DPVM Configuration Status*

```
switch# show dpvm status
DB is activated successfully, auto-learn is on
```

Example 11-2 *Displays the DPVM Current Dynamic Ports for the Specified VSAN*

```
switch# show dpvm ports vsan 10
-----
Interface Vsan Device pWWN                Device nWWN
-----
fc1/2      10    29:a0:00:05:30:00:6b:a0 fe:65:00:05:30:00:2b:a0
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 11-3 Displays the DPVM Config Database

```
switch# show dpvm database
pwwn 11:22:33:44:55:66:77:88 vsan 11
pwwn 22:22:33:44:55:66:77:88 vsan 22
pwwn 33:22:33:44:55:66:77:88 vsan 33
pwwn 44:22:33:44:55:66:77:88 vsan 44
[Total 4 entries]
```

Example 11-4 Displays the DPVM Active Database

```
switch# show dpvm database active
pwwn 11:22:33:44:55:66:77:88 vsan 22
pwwn 22:22:33:44:55:66:77:88 vsan 22
pwwn 33:22:33:44:55:66:77:88 vsan 33
[Total 3 entries]
* is auto-learnt entry
```

Example 11-5 Displays DPVM Config Database

```
switch# show dpvm database
pwwn 11:22:33:44:55:66:77:88 vsan 11
pwwn 22:22:33:44:55:66:77:88 vsan 22
pwwn 33:22:33:44:55:66:77:88 vsan 33
pwwn 44:22:33:44:55:66:77:88 vsan 44
[Total 4 entries]
```

Example 11-6 Compares Pending Database with the Config Database

```
switch# show dpvm pending-diff
Legend: "+" New Entry, "-" Missing Entry, "*" Possible Conflict Entry
-----
+ pwwn 55:22:33:44:55:66:77:88 vsan 55
- pwwn 11:22:33:44:55:66:77:88 vsan 11
* pwwn 44:22:33:44:55:66:77:88 vsan 44
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Sample DPVM Configuration

To configure a basic DPVM scenario, follow these steps:

- Step 1** Enable the DPVM feature and enable DPVM distribution.

```
switch1# config
Enter configuration commands, one per line. End with CNTL/Z.
switch1(config)# dpvm enable
switch1(config)# end
switch1# show dpvm database
switch1# show dpvm database active
switch1# show dpvm status
```

At this stage, the configuration does not have an active database and the **auto-learn** option is disabled.

- Step 2** Activate a null (empty) database so it can be populated with autolearned entries.

```
switch1# config
Enter configuration commands, one per line. End with CNTL/Z.
switch1(config)# dpvm activate
switch1(config)# dpvm commit
switch1(config)# end
switch1# show dpvm database
switch1# show dpvm database active
switch1# show dpvm status
```

At this stage, the database is successfully activated and the **auto-learn** option continues to be disabled.

- Step 3** Enable the auto-learn option and commit the configuration changes.

```
switch1# config
Enter configuration commands, one per line. End with CNTL/Z.
switch1(config)# dpvm auto-learn
switch1(config)# dpvm commit
switch1(config)# end
switch1# show dpvm database active
pwwn 21:00:00:e0:8b:0e:74:8a vsan 4(*)
pwwn 21:01:00:e0:8b:2e:87:8a vsan 5(*)
[Total 2 entries]
* is auto-learned entry
switch1# show dpvm ports
```

| Interface | Vsan | Device pWWN | Device nWWN |
|-----------|------|-------------------------|-------------------------|
| fc1/24 | 4 | 21:00:00:e0:8b:0e:74:8a | 20:00:00:e0:8b:0e:74:8a |
| fc1/27 | 5 | 21:01:00:e0:8b:2e:87:8a | 20:01:00:e0:8b:2e:87:8a |

```
switch1# show flogi database
```

| INTERFACE | VSAN | FCID | PORT NAME | NODE NAME |
|-----------|------|----------|-------------------------|-------------------------|
| fc1/24 | 4 | 0xe70100 | 21:00:00:e0:8b:0e:74:8a | 20:00:00:e0:8b:0e:74:8a |
| fc1/27 | 5 | 0xe80100 | 21:01:00:e0:8b:2e:87:8a | 20:01:00:e0:8b:2e:87:8a |

```
Total number of flogi = 2.

switch1# show dpvm status
DB is activated successfully, auto-learn is on
```

At this stage, the currently logged in devices (and their current VSAN assignment) populate the active database. However the entries are not yet permanent in the active database.

Send documentation comments to mdsfeedback-doc@cisco.com.

The output of the **show dpvm ports** and the **show flogi database** commands displays two other devices that have logged in (referred to as switch9 and switch3 in this sample configuration).

Step 4 Access switch9 and issue the following commands.

```
switch9# show dpvm database active
pwwn 21:00:00:e0:8b:0e:87:8a vsan 1(*)
pwwn 21:01:00:e0:8b:2e:74:8a vsan 1(*)
[Total 2 entries]
* is auto-learned entry
switch9# show dpvm status
DB is activated successfully, auto-learn is on
```

Step 5 Access switch3 and issue the following commands.

```
switch3# show dpvm database active
pwwn 21:00:00:e0:8b:0e:76:8a vsan 1(*)
pwwn 21:01:00:e0:8b:2e:76:8a vsan 1(*)
[Total 2 entries]
* is auto-learned entry
switch3# show dpvm status
DB is activated successfully, auto-learn is on
```

Step 6 Disable autolearning in switch1 and commit the configuration changes.

```
switch1# config
Enter configuration commands, one per line. End with CNTL/Z.
switch1(config)# no dpvm auto-learn
switch1(config)# dpvm commit
switch1(config)# end
switch1# show dpvm status
DB is activated successfully, auto-learn is off
switch1# show dpvm database active
pwwn 21:00:00:e0:8b:0e:74:8a vsan 4
pwwn 21:01:00:e0:8b:2e:87:8a vsan 5
pwwn 21:00:00:e0:8b:0e:87:8a vsan 1
pwwn 21:01:00:e0:8b:2e:74:8a vsan 1
pwwn 21:00:00:e0:8b:0e:76:8a vsan 1
pwwn 21:01:00:e0:8b:2e:76:8a vsan 1
[Total 6 entries]
* is auto-learned entry
switch1# show dpvm status
DB is activated successfully, auto-learn is off
```

At this stage, the autolearned entries are made permanent in the active database.

Send documentation comments to mdsfeedback-doc@cisco.com.

Step 7 Access switch9 and issue the following commands.

```
switch9# show dpvm database active
pwwn 21:00:00:e0:8b:0e:87:8a vsan 1
pwwn 21:01:00:e0:8b:2e:74:8a vsan 1
pwwn 21:00:00:e0:8b:0e:76:8a vsan 1
pwwn 21:01:00:e0:8b:2e:76:8a vsan 1
pwwn 21:00:00:e0:8b:0e:74:8a vsan 4
pwwn 21:01:00:e0:8b:2e:87:8a vsan 5
[Total 6 entries]
* is auto-learned entry
switch9# show dpvm status
DB is activated successfully, auto-learn is off
```

Step 8 Access switch3 and issue the following commands.

```
switch3# show dpvm database active
pwwn 21:00:00:e0:8b:0e:76:8a vsan 1
pwwn 21:01:00:e0:8b:2e:76:8a vsan 1
pwwn 21:00:00:e0:8b:0e:87:8a vsan 1
pwwn 21:01:00:e0:8b:2e:74:8a vsan 1
pwwn 21:00:00:e0:8b:0e:74:8a vsan 4
pwwn 21:01:00:e0:8b:2e:87:8a vsan 5
[Total 6 entries]
* is auto-learned entry
switch3# show dpvm status
DB is activated successfully, auto-learn is off
```



Note

These basic steps help you ascertain that the information is identical in all the switches in the fabric.

You have now configured a basic DPVM scenario in a Cisco MDS 9000 Family switch.

Default Settings

Table 11-1 lists the default settings for DPVM parameters.

Table 11-1 Default DPVM Parameters

| Parameters | Default |
|-------------------|-----------|
| DPVM | Disabled. |
| DPVM distribution | Enabled. |
| Autolearning | Disabled. |



Configuring Interfaces

A switch's main function is to relay frames from one data link to another. To do that, the characteristics of the interfaces through which the frames are received and sent must be defined. The configured interfaces can be Fibre Channel interfaces, the management interface (mgmt0), or VSAN interfaces.

This chapter describes the basic interface configuration to get your switch up and running. It includes the following sections:

- [Fibre Channel Interfaces, page 12-2](#)
- [Configuring Management Interfaces, page 12-19](#)
- [Configuring VSAN Interfaces, page 12-20](#)
- [Configuring CIM, page 12-20](#)
- [Displaying Interface Information, page 12-21](#)
- [Default Settings, page 12-33](#)



Note

See [Chapter 4, “Initial Configuration,”](#) and [Chapter 26, “Configuring IP Services,”](#) for more information on configuring mgmt0 interfaces.



Tip

Before you begin configuring the switch, ensure that the modules in the chassis are functioning as designed. To verify the status of a module at any time, issue the **show module** command in EXEC mode (see the [“Verifying the Module Status”](#) section on [page 4-15](#)).

Send documentation comments to mdsfeedback-doc@cisco.com.

Fibre Channel Interfaces

This section describes Fibre Channel interface characteristics, including (but not limited to) modes, states, and speeds. I includes the following sections:

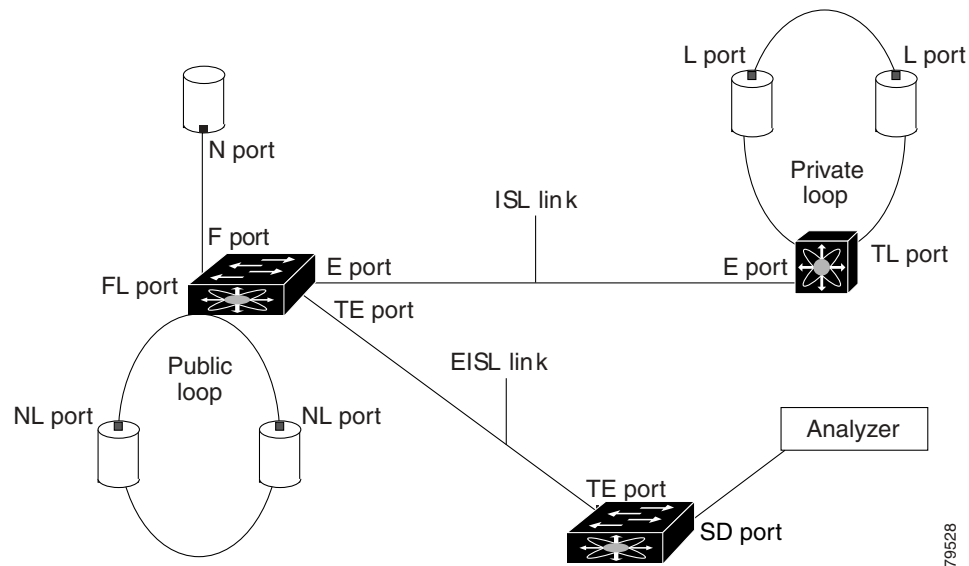
- [About Interface Modes, page 12-3](#)
- [About Interface States, page 12-6](#)
- [Configuring Fibre Channel Interface, page 12-9](#)
- [Graceful Shut Down, page 12-9](#)
- [Interface Modes, page 12-10](#)
- [TL Port ALPA Caches, page 12-10](#)
- [Administrative Speeds, page 12-11](#)
- [Interface Descriptions, page 12-12](#)
- [Buffer-to-Buffer Credits, page 12-12](#)
- [Performance Buffers, page 12-13](#)
- [Extended BB_credits, page 12-14](#)
- [Frame Encapsulation, page 12-16](#)
- [Receive Data Field Size, page 12-16](#)
- [Beacon Mode, page 12-16](#)
- [Identifying the Beacon LEDs, page 12-17](#)
- [Switch Port Attribute Default Values, page 12-17](#)
- [SFP Transmitter Types, page 12-18](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

About Interface Modes

Each physical Fibre Channel interface in a switch may operate in one of several port modes: E port, F port, FL port, TL port, TE port, SD port, ST port, and B port (see [Figure 12-1](#)). Besides these modes, each interface may be configured in auto or Fx port modes. These two modes determine the port type during interface initialization.

Figure 12-1 Cisco MDS 9000 Family Switch Port Modes



Note

Interfaces are created in VSAN 1 by default. See [Chapter 10, “Configuring and Managing VSANs.”](#)

Each interface has an associated administrative configuration and an operational status:

- The administrative configuration does not change unless you modify it. This configuration has various attributes that you can configure in administrative mode.
- The operational status represents the current status of a specified attribute like the interface speed. This status cannot be changed and is read-only. Some values may not be valid when the interface is down (for example, the operational speed).

A brief description of each interface mode follows.

E Port

In expansion port (E port) mode, an interface functions as a fabric expansion port. This port may be connected to another E port to create an Inter-Switch Link (ISL) between two switches. E ports carry frames between switches for configuration and fabric management. They serve as a conduit between switches for frames destined to remote N ports and NL ports. E ports support class 2, class 3, and class F service.

An E port connected to another switch may also be configured to form a PortChannel (see [Chapter 14, “Configuring PortChannels”](#)).

Send documentation comments to mdsfeedback-doc@cisco.com.

F Port

In fabric port (F port) mode, an interface functions as a fabric port. This port may be connected to a peripheral device (host or disk) operating as an N port. An F port can be attached to only one N port. F ports support class 2 and class 3 service.

FL Port

In fabric loop port (FL port) mode, an interface functions as a fabric loop port. This port may be connected to one or more NL ports (including FL ports in other switches) to form a public arbitrated loop. If more than one FL port is detected on the arbitrated loop during initialization, only one FL port becomes operational and the other FL ports enter nonparticipating mode. FL ports support class 2 and class 3 service.

TL Port

In translatable loop port (TL port) mode, an interface functions as a translatable loop port. It may be connected to one or more private loop devices (NL ports). TL ports are specific to Cisco MDS 9000 Family switches and have similar properties as FL ports. TL ports enable communication between a private loop device and one of the following devices:

- A device attached to any switch on the fabric
- A device on a public loop anywhere in the fabric
- A device on a different private loop anywhere in the fabric
- A device on the same private loop

TL ports support class 2 and class 3 services.

Private loop devices refer to legacy devices that reside on arbitrated loops. These devices are not aware of a switch fabric because they only communicate with devices on the same physical loop (see the [“Displaying TL Port Information”](#) section on page 12-30 and [“TL Port ALPA Caches”](#) section on page 12-10).



Tip

We recommend configuring devices attached to TL ports in zones that have up to 64 zone members.

TE Port

In trunking E port (TE port) mode, an interface functions as a trunking expansion port. It may be connected to another TE port to create an Enhanced ISL (EISL) between two switches. TE ports are specific to Cisco MDS 9000 Family switches. They expand the functionality of E ports to support the following:

- VSAN trunking
- Transport quality of service (QoS) parameters
- Fibre Channel trace (fctrace) feature

In TE port mode, all frames are transmitted in EISL frame format, which contains VSAN information. Interconnected switches use the VSAN ID to multiplex traffic from one or more VSANs across the same physical link. This feature is referred to as trunking in the Cisco MDS 9000 Family (see [Chapter 13, “Configuring Trunking”](#)). TE ports support class 2, class 3, and class F service.

Send documentation comments to mdsfeedback-doc@cisco.com.

SD Port

In SPAN destination port (SD port) mode, an interface functions as a switched port analyzer (SPAN). The SPAN feature is specific to switches in the Cisco MDS 9000 Family. It monitors network traffic that passes through a Fibre Channel interface. This monitoring is done using a standard Fibre Channel analyzer (or a similar switch probe) that is attached to an SD port. SD ports do not receive frames, they merely transmit a copy of the source traffic. The SPAN feature is nonintrusive and does not affect switching of network traffic for any SPAN source ports (see [Chapter 38, “Monitoring Network Traffic Using SPAN”](#)).

ST Port

In the SPAN tunnel port (ST port) mode, an interface functions as an entry point port in the source switch for the RSPAN Fibre Channel tunnel. The ST port mode and the remote SPAN (RSPAN) feature are specific to switches in the Cisco MDS 9000 Family. When configured in ST port mode, the interface cannot be attached to any device, and thus cannot be used for normal Fibre Channel traffic (see the [“Remote SPAN”](#) section on [page 38-15](#)).

Fx Port

Interfaces configured as Fx ports can operate in either F port or FL port mode. The Fx port mode is determined during interface initialization depending on the attached N port or NL port. This administrative configuration disallows interfaces to operate in any other mode—for example, preventing an interface to connect to another switch.

B Port

While E ports typically interconnect Fibre Channel switches, some SAN extender devices, such as the Cisco PA-FC-1G Fibre Channel port adapter, implement a bridge port (B port) model to connect geographically dispersed fabrics. This model uses B ports as described in the T11 Standard FC-BB-2.

[Figure 28-11 on page 28-34](#) depicts a typical SAN extension over an IP network.

If an FCIP peer is a SAN extender device that only supports Fibre Channel B ports, you need to enable the B port mode for the FCIP link. When a B port is enabled, the E port functionality is also enabled and they coexist. If the B port is disabled, the E port functionality remains enabled (see [Chapter 28, “Configuring IP Storage”](#)).

Auto Mode

Interfaces configured in auto mode can operate in one of the following modes: F port, FL port, E port, or TE port. The port mode is determined during interface initialization. For example, if the interface is connected to a node (host or disk), it operates in F port or FL port mode depending on the N port or NL port mode. If the interface is attached to a third-party switch, it operates in E port mode. If the interface is attached to another switch in the Cisco MDS 9000 Family, it may become operational in TE port mode (see [Chapter 13, “Configuring Trunking”](#)).

TL ports and SD ports are not determined during initialization and are administratively configured.

Send documentation comments to mdsfeedback-doc@cisco.com.

About Interface States

The interface state depends on the administrative configuration of the interface and the dynamic state of the physical link.

Administrative States

The administrative state refers to the administrative configuration of the interface as described in [Table 12-1](#).

Table 12-1 Administrative States

| Administrative State | Description |
|----------------------|---|
| Up | Interface is enabled. |
| Down | Interface is disabled. If you administratively disable an interface by shutting down that interface, the physical link layer state change is ignored. |

Operational States

The operational state indicates the current operational state of the interface as described in [Table 12-2](#).

Table 12-2 Operational States

| Operational State | Description |
|-------------------|--|
| Up | Interface is transmitting or receiving traffic as desired. To be in this state, an interface must be administratively up, the interface link layer state must be up, and the interface initialization must be completed. |
| Down | Interface cannot transmit or receive (data) traffic. |
| Trunking | Interface is operational in TE mode. |

Reason Codes

Reason codes are dependent on the operational state of the interface as described in [Table 12-3](#).

Table 12-3 Reason Codes for Interface States

| Administrative Configuration | Operational Status | Reason Code |
|------------------------------|--------------------|---|
| Up | Up | None. |
| Down | Down | Administratively down—If you administratively configure an interface as down, you disable the interface. No traffic is received or transmitted. |
| Up | Down | See Table 12-4 . |

Send documentation comments to mdsfeedback-doc@cisco.com.

If the administrative state is up and the operational state is down, the reason code differs based on the nonoperational reason code as described in [Table 12-4](#).

Table 12-4 Reason Codes for Nonoperational States

| Reason Code | Description | Applicable Modes |
|---|--|---------------------------|
| Link failure or not connected | The physical layer link is not operational. | All |
| SFP not present | The small form-factor pluggable (SFP) hardware is not plugged in. | |
| Initializing | The physical layer link is operational and the protocol initialization is in progress. | |
| Reconfigure fabric in progress | The fabric is currently being reconfigured. | |
| Offline | The Cisco SAN-OS software waits for the specified R_A_TOV time before retrying initialization. | |
| Inactive | The interface VSAN is deleted or is in a suspended state. To make the interface operational, assign that port to a configured and active VSAN. | |
| Hardware failure | A hardware failure is detected. | |
| Error disabled | Error conditions require administrative attention. Interfaces may be error-disabled for various reasons. For example: <ul style="list-style-type: none"> Configuration failure. Incompatible buffer-to-buffer credit configuration. To make the interface operational, you must first fix the error conditions causing this state; and next, administratively shut down or enable the interface. | |
| Isolation due to ELP failure | The port negotiation failed. | Only E ports and TE ports |
| Isolation due to ESC failure | The port negotiation failed. | |
| Isolation due to domain overlap | The Fibre Channel domains (fcdomain) overlap. | |
| Isolation due to domain ID assignment failure | The assigned domain ID is not valid. | |
| Isolation due to other side E port isolated | The E port at the other end of the link is isolated. | |
| Isolation due to invalid fabric reconfiguration | The port is isolated due to fabric reconfiguration. | |
| Isolation due to domain manager disabled | The fcdomain feature is disabled. | |
| Isolation due to zone merge failure | The zone merge operation failed. | |
| Isolation due to VSAN mismatch | The VSANs at both ends of an ISL are different. | |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 12-4 Reason Codes for Nonoperational States (continued)

| Reason Code | Description | Applicable Modes |
|---|--|-----------------------------|
| Nonparticipating | FL ports cannot participate in loop operations. It may happen if more than one FL port exists in the same loop, in which case all but one FL port in that loop automatically enters nonparticipating mode. | Only FL ports and TL ports |
| PortChannel administratively down | The interfaces belonging to the PortChannel are down. | Only PortChannel interfaces |
| Suspended due to incompatible speed | The interfaces belonging to the PortChannel have incompatible speeds. | |
| Suspended due to incompatible mode | The interfaces belonging to the PortChannel have incompatible modes. | |
| Suspended due to incompatible remote switch WWN | An improper connection is detected. All interfaces in a PortChannel must be connected to the same pair of switches. | |

32-Port Configuration Guidelines

The 32-port guidelines applies to the following hardware:

- The 32-port 2 Gbps or 1 Gbps switching module
- The Cisco MDS 9140 Switch

When configuring these host-optimized ports, the following port mode guidelines apply:

- You can configure only the first port in each 4-port group (for example, the first port in ports 1-4, the fifth port in ports 5-8 and so on) as an E port. If the first port in the group is configured as an E port, the other three ports in each group (ports 2-4, 6-8 and so on) are not usable and remain shutdown.
- If any of the other three ports are enabled, you cannot configure the first port as an E port. The other three ports continue to remain enabled.
- The auto mode is the default port mode. The auto mode is not allowed in a 32-port switching module or the host-optimized ports in the Cisco 9100 Series (16 host-optimized ports in the Cisco MDS 9120 switch and 32 host-optimized ports in the Cisco MDS 9140 switch).
- The default port mode is Fx (Fx negotiates to F or FL) for 32-port switching modules and the host-optimized ports in the Cisco 9100 Series (16 host-optimized ports in the Cisco MDS 9120 switch and 32 host-optimized ports in the Cisco MDS 9140 switch).
- The 32-port switching module does not support FICON.



Note

In the Cisco MDS 9100 Series, the left most groups of ports outlined in white (4 ports in the 9120 switch and 8 ports in the 9140 switch) are full line rate like the 16-port switching module. The other ports (16 ports in the 9120 switch and 32 ports in the 9140 switch) are host-optimized like the 32-port switching module. Each group of 4 host-optimized ports have the same rules as for the 32-port switching module.

Send documentation comments to mdsfeedback-doc@cisco.com.

Configuring Fibre Channel Interface

To configure a Fibre Channel interface, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/1 | Configures the specified interface. Note When a Fibre Channel interface is configured, it is automatically assigned a unique world wide name (WWN). If the interface's operational state is up, it is also assigned a Fibre Channel ID (FC ID). |

To configure a range of interfaces, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/1 - 4 , fc2/1 - 3 | Configures the range of specified interfaces. Note In this command, provide a space before and after the comma. |

Graceful Shut Down

Interfaces on a port are shut down by default (unless you modified the initial configuration).

As of Release 2.0(1b), the Cisco SAN-OS software implicitly performs a graceful shut down in response to either of the following actions for interfaces operating in the E port mode:

- If you shut down an interface
- If a Cisco SAN-OS software application executes a port shut down as part of its function

A graceful shut down ensures that no frames are lost when the interface is shutting down. When a shut down is triggered either by you or the Cisco SAN-OS software, the switches connected to the shut down link coordinate with each other to ensure that all frames in the ports are safely sent through the link before shutting down. This enhancement reduces the chance of frame loss.

A graceful shut down is not possible in the following situations:

- If you physically remove the port from the switch.
- If in-order-delivery (IOD) is enabled (see [“In-Order Delivery” section on page 24-11](#))
- If the Min_LS_interval interval is higher than 10 seconds (see [“Displaying Global FSPF Information” section on page 24-19](#))



Note

This feature is only triggered if both switches at either end of this E port interface are MDS switches and are using the Cisco SAN-OS Release 2.0(1b) (or later) software.

Send documentation comments to mdsfeedback-doc@cisco.com.

To gracefully shut down an interface, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/1 | Configures the specified interface. |
| Step 3 | switch(config-if)# shutdown | Gracefully shuts down the interface and administratively disables traffic flow (default). |

To enable traffic flow, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/1 | Configures the specified interface. |
| Step 3 | switch(config-if)# no shutdown | Enables traffic flow to administratively allow traffic when the no prefix is used (provided the operational state is up). |

Interface Modes

To configure the interface mode, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/1 switch(config-if)# | Configures the specified interface. |
| Step 3 | switch(config-if)# switchport mode F switch(config-if)# | Configures the administrative mode of the port. You can set the operational state to auto, E, F, FL, Fx, TL, or SD port mode. Note Fx ports refers to an F port or an FL port (host connection only), but not E ports. |
| | switch(config-if)# switchport mode auto switch(config-if)# | Configures the interface mode to auto-negotiate an E, F, FL, or TE port mode (not TL or SD port modes) of operation. Note TL ports and SD ports cannot be configured automatically. They must be administratively configured. |

TL Port ALPA Caches

Although TL ports cannot be automatically configured, as of Cisco SAN-OS Release 1.3(5) you can manually configure entries in arbitrated loop physical address (ALPA) caches. Generally, ALPA cache entries are automatically populated when an ALPA is assigned to a device. Each device is identified by its port world wide name (pWWN). When a device is allocated an ALPA, an entry for that device is automatically created in the ALPA cache.

A cache contains entries for recently allocated ALPA values. These caches are maintained on various TL ports. If a device already has an ALPA, the Cisco SAN-OS software attempts to allocate the same ALPA to the device each time. The ALPA cache is maintained in persistent storage and saves information across

Send documentation comments to mdsfeedback-doc@cisco.com.

switch reboots. The maximum cache size is 1000 entries. If the cache is full, and a new ALPA is allocated, the Cisco SAN-OS software discards an inactive cache entry (if available) to make space for the new entry.

See the “[Displaying TL Port Information](#)” section on page 12-30 for more information on TL ports.

To manually insert entries into the ALPA cache, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# tlport alpa-cache interface fc1/2 pwwn 22:00:00:20:37:46:09:bd alpa 0x02 | Configures manual entries into the ALPA cache. |
| Step 3 | switch(config)# tlport alpa-cache interface fc1/3 pwwn 22:00:00:20:37:46:09:bd | Removes this entry from the ALPA cache. |

Displaying the ALPA Cache Contents

The **show tlport alpa-cache** command displays the contents of the ALPA cache.

```
switch# show tlport alpa-cache
-----
alpa                pWWN                Interface
-----
0x02  22:00:00:20:37:46:09:bd  fc1/2
0x04  23:00:00:20:37:46:09:bd  fc1/2
```

The first entry indicates that if a device with a pWWN of 22:00:00:20:37:46:09:bd is exported on TL port fc1/2, then the pWWN is allocated an alpa 0x02 (if available).

Clearing the ALPA Cache

The **clear tlport alpa-cache** command clears the entire content of the ALPA cache.

Administrative Speeds

By default, the administrative speed for an interface is automatically calculated by the switch.



Caution

Changing the administrative speed is a disruptive operation.

To configure the administrative speed of the interface, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface mgmt0 | Selects the mgmt0 interface and enters interface configuration mode. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|---|--|
| Step 3 | <code>switch(config-if)# switchport speed 1000</code> <code>switch(config-if)#</code> | Configures the administrative speed of the interface to 1000 Mbps. The number indicates the speed in megabits per second (Mbps). You can set the speed to 1000 (for 1-Gbps interfaces), 2000 (for 2-Gbps interfaces), or auto (default). |
| | <code>switch(config-if)# switchport speed</code> <code>switch(config-if)#</code> | Reconfigures the factory default (auto) administrative speed of the interface. |

Interface Descriptions

To configure a description for an interface, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | <code>switch# config t</code> | Enters configuration mode. |
| Step 2 | <code>switch(config)# interface fc1/1</code> <code>switch(config-if)#</code> | Configures the specified interface. |
| Step 3 | <code>switch(config-if)# switchport description cisco-HBA2</code> | Configures the description of the interface. The string may be up to 80 characters long. |
| | <code>switch(config-if)# no switchport description</code> | Clears the description of the interface. |

Buffer-to-Buffer Credits

Buffer-to-buffer credits (BB_credits) are a flow control mechanism to ensure that FC switches do not run out of buffers, because switches must not drop frames. BB_credits are negotiated on a per-hop basis.

The receive BB_credit (`fcrxbbcredit`) value may be configured for each FC interface. In most cases, you do not need to modify the default configuration.



Note

The receive BB_credit values depend on the module type and the port mode. For 16-port switching modules and full rate ports, the default value is 16 for Fx mode and 255 for E or TE modes. The maximum value is 255 in all modes. This value can be changed as required. For 32-port switching modules and host-optimized ports, the default value is 12 for Fx, E, and TE modes. These values cannot be changed.



Note

In the Cisco MDS 9100 Series, the left most groups of ports outlined in white (4 ports in the 9120 switch and 8 ports in the 9140 switch) are full line rate like the 16-port switching module. The other ports (16 ports in the 9120 switch and 32 ports in the 9140 switch) are host-optimized like the 32-port switching module. Each group of 4 host-optimized ports have the same rules as for the 32-port switching module.

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure BB_credits for a Fibre Channel interface, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/1 switch(config-if)# | Configures the specified interface. |
| Step 3 | switch(config-if)# switchport fcrxbcredit default | Applies the default operational value to the selected interface. The operational value depends on the port mode. The default values are assigned based on the port capabilities. |
| | switch(config-if)# switchport fcrxbcredit 5 | Assigns a BB_credit of 5 to the selected interface. The range to assign BB_credits is between 1 and 255. |
| | switch(config-if)# switchport fcrxbcredit 5 mode E | Assigns this value if the port is operating in E or TE mode. The range to assign BB_credits is between 1 and 255. |
| | switch(config-if)# switchport fcrxbcredit 5 mode Fx | Assigns this value if the port is operating in F or FL mode. The range to assign BB_credits is between 1 and 255. |
| Step 4 | switch# do show int fc1/1 fc1/1 is up ... 16 receive B2B credit remaining 3 transmit B2B credit remaining | Displays the receive and transmit BB_credit along with other pertinent interface information for this interface. Note The BB_credit values are correct at the time the registers are read. They are useful to verify situations when the data traffic is slow. |

Performance Buffers

Regardless of the configured receive BB_credit value, additional buffers, called performance buffers, improve switch port performance. Instead of relying on the built-in switch algorithm, you can manually configure the performance buffer value for specific applications (for example, forwarding frames over FCIP interfaces).

For each physical Fibre Channel interface in any switch in the Cisco MDS 9000 Family, you can specify the amount of performance buffers allocated in addition to the configured receive BB_credit value.

The default performance buffer value is 0. If you use the **default** option, the built-in algorithm is used. If you do not specify this command, the **default** option is automatically used.

To configure performance buffers for a Fibre Channel interface, follow these steps:

| | Command | Purpose |
|--------|--|-------------------------------------|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/1 switch(config-if)# | Configures the specified interface. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|--|--|
| Step 3 | <code>switch(config-if)# switchport fcrxbbcredit performance-buffers 45</code> | Assigns a performance buffer of 45 to the selected interface. The value ranges from 1 and 145. |
| | <code>switch(config-if)# switchport fcrxbbcredit performance-buffers default</code> | Reverts to the factory default of using the built-in algorithm. |



Note Use the **show interface bbcredit** command to display performance buffer values and other BB_credit information.


Extended BB_credits

The BB_credits feature allows you to configure up to 255 receive buffers. This number is insufficient for long haul links. To facilitate BB_credits for long haul links, the extended BB_credits flow control mechanism introduced in Cisco SAN-OS Release 2.0(1b). This feature allows you to configure up to 3,500 receive BB_credits on a Fibre Channel port.

To use this feature, you must meet the following requirements:

- Obtain the ENTERPRISE_PKG license (see [Chapter 3, “Obtaining and Installing Licenses”](#)).
- Configure this feature in any port of the full-rate 4-port group in either the Cisco MDS 9216i Switch or in the MPS-14/2 module (see [Figure 12-1](#)).

Figure 12-2 Port Group Support for the Extended BB_Credits Feature



The port groups that support extended credits configurations are as follows.

- Any one port in ports 1to 4 (identified as Group 1 in [Figure 12-1](#)).
- Any one port in ports 5 to 8 (identified as Group 2 in [Figure 12-1](#)).
- Any one port in ports 9 to 12 (identified as Group 3 in [Figure 12-1](#)).



Note The last two Fibre Channel ports (Port 13 and Port 14) and the two Gigabit Ethernet ports do not support the extended BB_credits feature (see [Figure 12-1](#)).

- Explicitly enable this feature in the required Cisco MDS switch.

Send documentation comments to mdsfeedback-doc@cisco.com.

- Disable the remaining three ports in the 4-port group if you need to assign more than 2,400 BB_credits to the first port in the port group.
 - If you assign less than 2,400 extended BB_credits to any one port in a port group, the remaining three ports in that port group can retain up to 255 BB_credits based on the port mode.


Note

The receive BB_credit value for the remaining three ports depends on the port mode. The default value is 16 for the Fx mode and 255 for E or TE modes. The maximum value is 255 in all modes. This value can be changed as required without exceeding the maximum value of 255 BB_credits.

- If you assign more than 2,400 (up to a maximum of 3,500) extended BB_credits to the port in a port group, you must disable the other three ports.
- Be aware that changing the BB_credits value results in the port being disabled and then reenabled.
- Disable (explicitly) this feature if you need to nondisruptive downgrade to Cisco SAN-OS Release 1.3 or earlier. When you disable this feature, the existing extended BB_credit configuration is completely erased.


Note

The extended BB_credit configuration takes precedence over the receive BB_credit and performance buffer configurations.

To configure extended BB_credits for a MDS-14/2 interface or for an interface in a Cisco MDS 9216i switch, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fcrxbbcredit extended enable switch(config)# no fcrxbbcredit extended enable | Enables the extended BB_credits feature. Disables (default) the extended BB_credits feature. |
| Step 3 | switch(config)# interface fc1/1 switch(config-if)# | Configures the specified interface. |
| Step 4 | switch(config-if)# switchport fcrxbbcredit extended 1500 switch(config-if)# no switchport fcrxbbcredit extended 2500 | Applies the extended BB_credit value of 1,500 credits to the selected interface. The valid range is from 256 to 3,500 credits. Clears the configured BB_credit configuration for this port. |
| Step 5 | switch# do show int fc3/2 fc3/2 is trunking Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN) Port WWN is 20:82:00:05:30:00:2a:1e Peer port WWN is 20:42:00:0b:46:79:f1:80 Admin port mode is auto, trunk mode is on Port mode is TE Port vsan is 1 Speed is 2 Gbps Transmit B2B Credit is 255 Receive B2B Credit is 1500 Receive data field Size is 2112 ... | Displays the receive and transmit BB_credit values along with other pertinent interface information for this interface if the interface is in the up state. Note The receive BB_credit value reflects the extended BB_credit configuration, if applicable. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Frame Encapsulation

The **switchport encap eisl** command only applies to SD port interfaces. This command determines the frame format for all frames transmitted by the interface in SD port mode. If the encap is set to EISL, all frames are transmitted in the EISL frame format irrespective of the SPAN source(s).

The **switchport encap eisl** command is disabled by default. If you enable encapsulation, all outgoing frames are encapsulated, and you will see a new line (*Encapsulation is eisl*) in the **show interface SD_port_interface** command output (see the “[Encapsulating Frames](#)” section on page 38-8).

Receive Data Field Size

You can also configure the receive data field size for Fibre Channel interfaces. If the default data field size is 2112 bytes, the frame length will be 2148 bytes.

Use the **switchport fcrxbufsize** command to configure the data field size for Fibre Channel interfaces.

To configure data field size, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/1 switch(config-if)# | Configures the specified interface. |
| Step 3 | switch(config-if)# switchport fcrxbufsize 2000 | Reduces the data field size for the selected interface to 2000 bytes. The default is 2112 bytes and the range is from 256 to 2112 bytes. |

Beacon Mode

By default, the beacon mode is disabled on all switches. The beacon mode is indicated by a flashing green light that helps you identify the physical location of the specified interface.

The **beacon** command has no effect on the operation of the interface.

To enable beacon mode for a specified interface or range of interfaces, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/1 switch(config-if)# | Configures the specified interface. |
| Step 3 | switch(config-if)# switchport beacon switch(config-if)# no switchport beacon | Enables the beacon mode for the interface. Disables the beacon mode for the interface. |



Note


The flashing green light turns on automatically when an external loopback is detected that causes the interfaces to be isolated. The flashing green light overrides the beacon mode configuration. The state of the LED is restored to reflect the beacon mode configuration after the external loopback is removed.

Send documentation comments to mdsfeedback-doc@cisco.com.

Identifying the Beacon LEDs

Figure 12-3 displays the status, link, and speed LEDs in a 16-port switching module.

Figure 12-3 Cisco MDS 9000 Family Switch Interface Modes



| | | | |
|----------|--|----------|--|
| 1 | Status LED ¹ | 3 | Link LEDs ¹ and speed LEDs ² |
| 2 | 1/2-Gbps Fibre Channel port group ³ | 4 | Asset tag ⁴ |

1. See the “Identifying Module LEDs” section on page 7-9.
2. See the “About Speed LEDs” section on page 12-17.
3. See the “32-Port Configuration Guidelines” section on page 12-8.
4. Refer to the *Cisco MDS 9000 Family Hardware Installation Guide*.

About Speed LEDs

Each port has one link LED on the left and one speed LED on the right.

The speed LED displays the speed of the port interface:

- Off—The interface attached to that port is functioning at 1000 Mbps.
- On (solid green)—The interface attached to that port is functioning at 2000 Mbps (for 2 Gbps interfaces).

The speed LED also displays if the beacon mode is enabled or disabled:

- Off—Beacon mode is disabled.
- On (flashing green)—The beacon mode is enabled. The LED flashes at one-second intervals.

Switch Port Attribute Default Values

You can configure attribute default values for various switch port attributes. These attributes will be applied globally to all future switch port configurations, even if you do not individually specify them at that time.

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure switch port attributes, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# no system default switchport shutdown switch(config-if)# | Configures the default setting for administrative state of an interface as Up. (The factory default setting is Down). Tip This command is applicable only to interfaces for which no user configuration exists for the administrative state. |
| | switch(config)# system default switchport shutdown switch(config-if)# | Configures the default setting for administrative state of an interface as Down. This is the factory default setting. Tip This command is applicable only to interfaces for which no user configuration exists for the administrative state. |
| | switch(config)# system default switchport trunk mode auto switch(config-if)# | Configures the default setting for administrative trunk mode state of an interface as Auto. (The factory default setting is trunk mode On). |

SFP Transmitter Types

As of Release 2.0(1b), the term FCOT, Fibre Channel optical transmitter, is replaced by the term SFP, small form-factor pluggable, in the Cisco SAN-OS software.

The SFP hardware transmitters are identified by their acronyms when displayed in the **show interface brief** command. If the related SFP has a Cisco-assigned extended ID, then the **show interface** and **show interface brief** commands display the ID instead of the transmitter type. The **show interface transceiver** command and the **show interface fcslot/port transceiver** command display both values for Cisco supported SFPs. Table 12-5 defines the acronyms used in the command output (see the “Displaying Interface Information” section on page 12-21).

Table 12-5 SFP Transmitter Acronym Definitions

| Definition | Acronym |
|---|---------|
| Standard transmitters defined in the GBIC specifications | |
| short wave laser | swl |
| long wave laser | lwl |
| long wave laser cost reduced | lwcr |
| electrical | elec |
| Extended transmitters assigned to Cisco-supported SFPs | |
| CWDM-1470 | c1470 |
| CWDM-1490 | c1490 |
| CWDM-1510 | c1510 |
| CWDM-1530 | c1530 |
| CWDM-1550 | c1550 |
| CWDM-1570 | c1570 |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 12-5 SFP Transmitter Acronym Definitions (continued)

| Definition | Acronym |
|---|---------|
| Standard transmitters defined in the GBIC specifications | |
| CWDM-1590 | c1590 |
| CWDM-1610 | c1610 |

Configuring Management Interfaces

You can remotely configure the switch through the management interface (mgmt0). To configure a connection remotely, you must configure the IP parameters (IP address, subnet mask, and default gateway) from the CLI so that the switch is reachable.



Note

Before you begin to configure the management interface manually, obtain the switch's IP address and IP subnet mask.

To configure the mgmt0 Ethernet interface, follow these steps:

| | Command | Purpose |
|---------------|--|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# interface mgmt0 switch(config-if)# | Configures the management Ethernet interface on the switch to configure the management interface. |
| Step 3 | switch(config-if)# ip address 172.16.1.2 255 255.255.0 | Enters the IP address and IP subnet mask for the interface specified in Step 2. |
| Step 4 | switch(config-if)# no shutdown | Enables the interface. |
| Step 5 | switch(config-if)# exit switch(config)# | Returns to configuration mode. |
| Step 6 | switch(config)# ip default-gateway 1.1.1.4 switch(config)# | Configures the default gateway IP address. |
| Step 7 | switch(config)# exit switch# | Returns to EXEC mode. |
| Step 8 | switch# copy running-config startup-config | (Optional) Saves your configuration changes to the file system. Note If you wish to save your configuration, you can issue this command at any time. |

The management port (mgmt0) is autosensing and operates in full duplex mode at a speed of 10/100 Mbps. The speed and mode cannot be configured.



Note

You need to explicitly configure a default gateway to connect to the switch and send IP packets or add a route for each subnet.

Send documentation comments to mdsfeedback-doc@cisco.com.

Configuring VSAN Interfaces

VSANs apply to Fibre Channel fabrics and enable you to configure multiple isolated SAN topologies within the same physical infrastructure. You can create an IP interface on top of a VSAN and then use this interface to send frames to this VSAN. To use this feature, you must configure the IP address for this VSAN. VSAN interfaces cannot be created for nonexistent VSANs.

Follow these guidelines when creating or deleting VSAN interfaces:

- Create a VSAN before creating the interface for that VSAN. If a VSAN does not exist, the interface cannot be created.
- Create the interface VSAN—it is not created automatically.
- If you delete the VSAN, the attached interface is automatically deleted.
- Configure each interface only in one VSAN.



Tip

After configuring the VSAN interface, you can configure an IP address or Virtual Router Redundancy Protocol (VRRP) features (see [Chapter 26, “Configuring IP Services”](#)).

To create a VSAN interface, follow these steps:

| | Command | Purpose |
|--------|---|----------------------------------|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface vsan 5 switch(config-if)# | Configures a VSAN with the ID 5. |

Configuring CIM

Common Information Model (CIM) is an object-oriented information model that extends the existing standards for describing management information in a network/enterprise environment. CIM messages are independent of platform and implementation because they are encoded in N Extensible Markup Language (XML). CIM consists of a specification and a schema. The specification defines the syntax and rules for describing management data and integrating with other management models. The schema provides the actual model descriptions for systems, applications, networks, and devices.

For more information about CIM, refer to the specification available through the Distributed Management Task Force (DMTF) website at the following URL: <http://www.dmtf.org/>

For further information about Cisco MDS 9000 Family support for CIM servers, refer to the *Cisco MDS 9000 Family CIM Programming Reference Guide*.

A CIM client is required to access the CIM server. The client can be any client that supports CIM.

Added Security on a CIM Server

For added security, you can install an SSL certificate to encrypt the logon information and enable the HTTPS server before enabling the CIM server. The CIM server is disabled by default. If you do not enable the HTTPS server, the standard HTTP server is enabled (default).

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure a CIM server using the HTTPS protocol, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# cimserver certificate bootflash:simserver.pem | Installs a Secure Socket Layer (SSL) certificate specified in the file named with a .pem extension. |
| | switch(config)# cimserver clearcertificate Certificate1 | Optional. Clears the specified SSL certificate (Certificate1). |
| Step 3 | switch(config)# cimserver enableHttps | Enables HTTPS (secure protocol). |
| | switch(config)# no cimserver enableHttps | Optional. Disables HTTPS (default). |
| Step 4 | switch(config)# cimserver enable | Enables the CIM server. |
| | switch(config)# no cimserver enable | Optional. Disables the CIM server (default). |

To configure a CIM server using the HTTP protocol, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# cimserver enable | Enables the CIM server using the default HTTP (non-secure) protocol. |
| | switch(config)# no cimserver enable | Optional. Disables the CIM server (default). |
| | switch(config)# no cimserver enableHttp | Optional. Disables HTTP. |
| | switch(config)# cimserver enableHttp | Optional. Enables HTTP and reverts to the switch default. |

Displaying Interface Information

The **show interface** command is invoked from the EXEC mode and displays the interface configurations. Without any arguments, this command displays the information for all the configured interfaces in the switch. See Examples 12-1 to 12-15.

Example 12-1 Displays All Interfaces

```
switch# show interface
fc1/1 is up
  Hardware is Fibre Channel, SFP is short wave laser
  Port WWN is 20:0b:00:05:30:00:8d:de
  Admin port mode is F
  Port mode is F, FCID is 0x610000
  Port vsan is 2
  Speed is 2 Gbps
  Transmit B2B Credit is 3
  Receive B2B Credit is 16
  Receive data field Size is 2112
  Beacon is turned off
  5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    134 frames input, 8468 bytes
      0 discards, 0 errors
      0 CRC, 0 unknown class
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```

    0 too long, 0 too short
    154 frames output, 46072 bytes
    0 discards, 0 errors
    1 input OLS, 1 LRR, 0 NOS, 0 loop inits
    1 output OLS, 0 LRR, 1 NOS, 0 loop inits
    16 receive B2B credit remaining
    3 transmit B2B credit remaining.
. . .
fc1/9 is trunking
    Hardware is Fibre Channel, SFP is long wave laser cost reduced
    Port WWN is 20:09:00:05:30:00:97:9e
    Peer port WWN is 20:0b:00:0b:5f:a3:cc:00
    Admin port mode is E, trunk mode is on
    Port mode is TE
    Port vsan is 100
    Speed is 2 Gbps
    Transmit B2B Credit is 255
    Receive B2B Credit is 255
    Receive data field Size is 2112
    Beacon is turned off
    Trunk vsans (admin allowed and active) (1,100,3000)
    Trunk vsans (up) (1,100,3000)
    Trunk vsans (isolated) ()
    Trunk vsans (initializing) ()
    5 minutes input rate 280 bits/sec, 35 bytes/sec, 0 frames/sec
    5 minutes output rate 176 bits/sec, 22 bytes/sec, 0 frames/sec
    4609939 frames input, 8149405708 bytes
        0 discards, 0 errors
        0 CRC, 0 unknown class
        0 too long, 0 too short
    4638491 frames output, 7264731728 bytes
        0 discards, 0 errors
        3 input OLS, 9 LRR, 1 NOS, 0 loop inits
        9 output OLS, 7 LRR, 1 NOS, 0 loop inits
        16 receive B2B credit remaining
        3 transmit B2B credit remaining.
. . .
fc1/13 is up
    Hardware is Fibre Channel, SFP is short wave laser
    Port WWN is 20:0d:00:05:30:00:97:9e
    Admin port mode is auto, trunk mode is on
    Port mode is F, FCID is 0x650100
    Port vsan is 100
    Speed is 2 Gbps
    Transmit B2B Credit is 3
    Receive B2B Credit is 16
    Receive data field Size is 2112
    Beacon is turned off
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    8696 frames input, 3227212 bytes
        0 discards, 0 errors
        0 CRC, 0 unknown class
        0 too long, 0 too short
    16799 frames output, 6782444 bytes
        0 discards, 0 errors
        0 input OLS, 0 LRR, 0 NOS, 0 loop inits
        1 output OLS, 1 LRR, 0 NOS, 1 loop inits
        16 receive B2B credit remaining
        3 transmit B2B credit remaining.
. . .

```


Send documentation comments to mdsfeedback-doc@cisco.com.

```

sup-fc0 is up
  Hardware is Fibre Channel
  Speed is 1 Gbps
  139597 packets input, 13852970 bytes
    0 multicast frames, 0 compressed
    0 input errors, 0 frame, 0 overrun 0 fifo
  139516 packets output, 16759004 bytes, 0 underruns
    0 output errors, 0 collisions, 0 fifo
    0 carrier errors

mgmt0 is up
  Hardware is FastEthernet
  Address is 0005.3000.80fe
  Internet address is 172.19.48.96/25
  MTU 1500 bytes, BW 100 Mbps
  321561 packets input, 70215667 bytes
    0 multicast frames, 0 compressed
    0 input errors, 0 frame, 0 overrun 0 fifo
  334550 packets output, 307482596 bytes, 0 underruns
    0 output errors, 0 collisions, 0 fifo
    0 carrier errors

vsan1 is up, line protocol is up
  WWPN is 10:00:00:05:30:00:12:63, FCID is 0xef001e
  Internet address is 10.10.11.10/24
  MTU 1500 bytes, BW 1000000 Kbit
  0 packets input, 0 bytes, 0 errors, 0 multicast
  0 packets output, 0 bytes, 0 errors, 0 dropped
. . .
port-channel 1 is trunking
  Hardware is Fibre Channel
  Port WWN is 24:01:00:05:30:00:97:9e
  Admin port mode is E, trunk mode is on
  Port mode is TE
  Port vsan is 1
  Speed is 4 Gbps
  Trunk vsans (admin allowed and active) (1,100,3000)
  Trunk vsans (up) (1)
  Trunk vsans (isolated) (100,3000)
  Trunk vsans (initializing) ( )
  5 minutes input rate 648 bits/sec, 81 bytes/sec, 0 frames/sec
  5 minutes output rate 304 bits/sec, 38 bytes/sec, 0 frames/sec
  4629945 frames input, 206672020 bytes
    0 discards, 0 errors
    0 CRC, 0 unknown class
    0 too long, 0 too short
  4547515 frames output, 687414748 bytes
    0 discards, 0 errors
    2 input OLS, 2 LRR, 4 NOS, 0 loop inits
    6 output OLS, 2 LRR, 4 NOS, 0 loop inits
  Member[1] : fc1/1
  Member[2] : fc1/2.
. . .

```

You can also specify arguments (a range of interfaces or multiple, specified interfaces) to display interface information. You can specify a range of interfaces by issuing a command with the following example format:

interface fc1/1 - 5 , fc2/5 - 7



Note

The spaces are required before and after the dash (-) and before and after the comma (,).

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 12-2 Displays Multiple, Specified Interfaces

```
switch# show interface fc3/13 , fc3/16
fc3/13 is up
  Hardware is Fibre Channel, SFP is short wave laser
  Port WWN is 20:8d:00:05:30:00:97:9e
  Admin port mode is FX
  Port mode is F, FCID is 0x7b0300
  Port vsan is 1
  Speed is 2 Gbps
  Transmit B2B Credit is 3
  Receive B2B Credit is 12
  Receive data field Size is 2112
  Beacon is turned off
  5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    1856 frames input, 116632 bytes
      0 discards, 0 errors
      0 CRC, 0 unknown class
      0 too long, 0 too short
    1886 frames output, 887712 bytes
      0 discards, 0 errors
      0 input OLS, 0 LRR, 0 NOS, 1 loop inits
      1 output OLS, 1 LRR, 0 NOS, 1 loop inits
      16 receive B2B credit remaining
      3 transmit B2B credit remaining.

fc3/16 is up
  Hardware is Fibre Channel, SFP is short wave laser
  Port WWN is 20:90:00:05:30:00:97:9e
  Admin port mode is FX
  Port mode is F, FCID is 0x7d0100
  Port vsan is 3000
  Speed is 2 Gbps
  Transmit B2B Credit is 3
  Receive B2B Credit is 12
  Receive data field Size is 2112
  Beacon is turned off
  5 minutes input rate 504 bits/sec, 63 bytes/sec, 0 frames/sec
  5 minutes output rate 520 bits/sec, 65 bytes/sec, 0 frames/sec
    47050 frames input, 10311824 bytes
      0 discards, 0 errors
      0 CRC, 0 unknown class
      0 too long, 0 too short
    62659 frames output, 10676988 bytes
      0 discards, 0 errors
      0 input OLS, 0 LRR, 0 NOS, 0 loop inits
      1 output OLS, 1 LRR, 0 NOS, 1 loop inits
      16 receive B2B credit remaining
      3 transmit B2B credit remaining.
```

Example 12-3 Displays a Specific Interface

```
switch# show interface fc2/2
fc2/2 is trunking
  Port description is Trunk to Core-4
  Hardware is Fibre Channel, SFP is short wave laser
  Port WWN is 20:42:00:05:30:00:97:9e
  Peer port WWN is 20:cc:00:05:30:00:50:9e
  Admin port mode is E, trunk mode is on
  Port mode is TE
  Port vsan is 1
  Speed is 2 Gbps
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
Transmit B2B Credit is 255
Receive B2B Credit is 255
Receive data field Size is 2112
Beacon is turned off
Belongs to port-channel 2
Trunk vsans (admin allowed and active) (1,100,3000)
Trunk vsans (up) (1)
Trunk vsans (isolated) (100,3000)
Trunk vsans (initializing) ()
5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
5 minutes output rate 32 bits/sec, 4 bytes/sec, 0 frames/sec
  2214834 frames input, 98673588 bytes
    0 discards, 0 errors
    0 CRC, 0 unknown class
    0 too long, 0 too short
  2262415 frames output, 343158368 bytes
    0 discards, 0 errors
  1 input OLS, 1 LRR, 1 NOS, 0 loop inits
  2 output OLS, 1 LRR, 0 NOS, 0 loop inits
  16 receive B2B credit remaining
  3 transmit B2B credit remaining.
```

Example 12-4 Displays a VSAN Interface

```
switch# show interface vsan 2
vsan2 is up, line protocol is up
  WWPN is 10:00:00:05:30:00:59:1f, FCID is 0xb90100
  Internet address is 10.1.1.1/24
  MTU 1500 bytes, BW 1000000 Kbit
  0 packets input, 0 bytes, 0 errors, 0 multicast
  0 packets output, 0 bytes, 0 errors, 0 dropped
```

Example 12-5 Displays CIM Server Certificate Files

```
switch# show cimserver certificateName
cimserver certificate file name is servcert.pem
```

Example 12-6 Displays the CIM Server Configuration

```
switch# show cimserver
cimserver is enabled
cimserver Http is not enabled
cimserver Https is enabled
cimserver certificate file name is servcert.pem
```

Example 12-7 Displays the CIM Server HTTPS Status

```
switch# show cimserver httpsstatus
cimserver Https is enabled
```

Example 12-8 Displays the CIM Server HTTP Status

```
switch# show cimserver httpstatus
cimserver Http is not enabled
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 12-9 Displays Port Description

```
switch# show interface description
```

```
-----
Interface      Description
-----
fc3/1          test intest
fc3/2          --
fc3/3          --
fc3/4          TE port
fc3/5          --
fc3/6          --
fc3/10         Next hop switch 5
fc3/11         --
fc3/12         --
fc3/16         --
-----
```

```
-----
Interface      Description
-----
port-channel 1  --
port-channel 5  --
port-channel 6  --
-----
```

Example 12-10 Displays Interface Information in a Brief Format

```
switch# show interface brief
```

```
-----
Interface  Vsan  Admin  Admin  Status      SFP  Oper  Oper  Port
          Mode  Trunk  Mode
          Mode
-----
fc1/1      1      E      on     trunking    swl  TE    2    1
fc1/2      1      E      on     trunking    swl  TE    2    1
fc1/3      1      auto   on     SFPAbsent   --   --    --   --
fc1/4      1      auto   on     SFPAbsent   --   --    --   --
fc1/5      3000   auto   on     up          swl  F     2    --
...
fc2/2      1      E      on     trunking    swl  TE    2    2
fc2/3      1      auto   on     down        c1610 --    --   --
fc2/4      1      auto   on     down        c1590 --    --   --
fc2/5      3000   auto   on     notConnected lwcr --    --   --
fc2/6      1      auto   on     SFPAbsent   --   --    --   --
...
fc3/16     3000   FX     --     up          swl  F     2    --
fc3/17     1      FX     --     SFPAbsent   --   --    --   --
-----
```

```
-----
Interface      Status      IP Address      Speed      MTU
-----
GigabitEthernet4/1  SFPAbsent  --             auto       1500
...
GigabitEthernet4/6  down       10.1.1.2/8     auto       3000
GigabitEthernet4/7  down       10.1.1.27/24   auto       1500
GigabitEthernet4/8  down       --             auto       1500
-----
```

```
-----
Interface      Status      Oper Mode      Oper Speed
                  (Gbps)
-----
```

```
iscsi4/1      down       --
...
-----
```

Send documentation comments to mdsfeedback-doc@cisco.com.

| Interface | Status | Speed (Gbps) | | | |
|-----------|--------|-----------------|--|--|--|
| sup-fc0 | up | 1 | | | |

| Interface | Status | IP Address | Speed | MTU |
|-----------|--------|-----------------|----------|------|
| mgmt0 | up | 172.19.48.96/25 | 100 Mbps | 1500 |

| Interface | Vsan | Admin Trunk Mode | Status | Oper Mode | Oper Speed (Gbps) |
|----------------|------|------------------------|----------|--------------|-------------------------|
| port-channel 1 | 1 | on | trunking | TE | 4 |
| port-channel 2 | 1 | on | trunking | TE | 4 |

| Interface | Vsan | Admin Mode | Admin Trunk Mode | Status | Oper Mode | Profile | Port-channel |
|-----------|------|---------------|------------------------|--------------|--------------|---------|--------------|
| fcip10 | 1 | auto | on | notConnected | -- | 10 | -- |

Example 12-11 Displays Interface Counters

```
switch# show interface counters
fc3/1
  5 minutes input rate 24 bits/sec, 3 bytes/sec, 0 frames/sec
  5 minutes output rate 16 bits/sec, 2 bytes/sec, 0 frames/sec
  3502 frames input, 268400 bytes
    0 discards, 0 CRC, 0 unknown class
    0 too long, 0 too short
  3505 frames output, 198888 bytes
    0 discards
  1 input OLS, 1 LRR, 1 NOS, 0 loop inits
  2 output OLS, 1 LRR, 1 NOS, 0 loop inits
  1 link failures, 1 sync losses, 1 signal losses
.
.
.
fc9/8
  5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  0 frames input, 0 bytes
    0 class-2 frames, 0 bytes
    0 class-3 frames, 0 bytes
    0 class-f frames, 0 bytes
    0 discards, 0 CRC, 0 unknown class
    0 too long, 0 too short
  0 frames output, 0 bytes
    0 class-2 frames, 0 bytes
    0 class-3 frames, 0 bytes
    0 class-f frames, 0 bytes
    0 discards
  0 input OLS, 0 LRR, 0 NOS, 0 loop inits
  0 output OLS, 0 LRR, 0 NOS, 0 loop inits
  0 link failures, 0 sync losses, 0 signal losses
  16 receive B2B credit remaining
  3 transmit B2B credit remaining.
. . .
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```

sup-fc0
  114000 packets input, 11585632 bytes
    0 multicast frames, 0 compressed
    0 input errors, 0 frame, 0 overrun 0 fifo
  113997 packets output, 10969672 bytes, 0 underruns
    0 output errors, 0 collisions, 0 fifo
    0 carrier errors

mgmt0
  31557 packets input, 2230860 bytes
    0 multicast frames, 0 compressed
    0 input errors, 0 frame, 0 overrun 0 fifo
  26618 packets output, 16824342 bytes, 0 underruns
    0 output errors, 0 collisions, 7 fifo
    0 carrier errors

vsan1
  0 packets input, 0 bytes, 0 errors, 0 multicast
  0 packets output, 0 bytes, 0 errors, 0 dropped
.
.
.
port-channel 1
  5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  0 frames input, 0 bytes
    0 class-2 frames, 0 bytes
    0 class-3 frames, 0 bytes
    0 class-f frames, 0 bytes
    0 discards, 0 CRC, 0 unknown class
    0 too long, 0 too short
  0 frames output, 0 bytes
    0 class-2 frames, 0 bytes
    0 class-3 frames, 0 bytes
    0 class-f frames, 0 bytes
    0 discards
  0 input OLS, 0 LRR, 0 NOS, 0 loop inits
  0 output OLS, 0 LRR, 0 NOS, 0 loop inits
  0 link failures, 0 sync losses, 0 signal losses

```



Note Interfaces 9/8 and 9/9 are not trunking ports and display class 2, 3, and F information as well.

Example 12-12 Displays Interface Counters in Brief Format

```
switch# show interface counters brief
```

| Interface | Input (rate is 5 min avg) | | Output (rate is 5 min avg) | |
|-----------|---------------------------|-----------------|----------------------------|-----------------|
| | Rate Mbits/s | Total Frames | Rate Mbits/s | Total Frames |
| fc3/1 | 0 | 3871 | 0 | 3874 |
| fc3/2 | 0 | 3902 | 0 | 4232 |
| fc3/3 | 0 | 3901 | 0 | 4138 |
| fc3/4 | 0 | 3895 | 0 | 3894 |
| fc3/5 | 0 | 3890 | 0 | 3897 |
| fc9/8 | 0 | 0 | 0 | 0 |
| fc9/9 | 0 | 5 | 0 | 4 |
| fc9/10 | 0 | 4186 | 0 | 4182 |
| fc9/11 | 0 | 4331 | 0 | 4315 |

Send documentation comments to mdsfeedback-doc@cisco.com.

| Interface | Input (rate is 5 min avg) | | Output (rate is 5 min avg) | |
|----------------|---------------------------|--------|----------------------------|--------|
| | Rate | Total | Rate | Total |
| | Mbits/s | Frames | Mbits/s | Frames |
| port-channel 1 | 0 | 0 | 0 | 0 |
| port-channel 2 | 0 | 3946 | 0 | 3946 |

Example 12-13 Displays BB_credit Information

```
switch# show interface bbbcredit
fc2/1 is down (SFP not present)
...
fc2/17 is trunking
    Transmit B2B Credit is 255
    Receive B2B Credit is 12
    Receive B2B Credit performance buffers is 375
        12 receive B2B credit remaining
        255 transmit B2B credit remaining
fc2/18 is down (SFP not present)
fc2/19 is down (SFP not present)
fc2/20 is down (SFP not present)
fc2/21 is down (Link failure or not-connected)
...
fc2/31 is up
    Transmit B2B Credit is 0
    Receive B2B Credit is 12
    Receive B2B Credit performance buffers is 48
        12 receive B2B credit remaining
        0 transmit B2B credit remaining
fc2/32 is down (Link failure or not-connected)
```

Example 12-14 Displays BB_credit Information for a Specified Fibre Channel Interface

```
switch# show interface fc2/31 bbbcredit
fc2/31 is up
    Transmit B2B Credit is 0
    Receive B2B Credit is 12
    Receive B2B Credit performance buffers is 48
        12 receive B2B credit remaining
        0 transmit B2B credit remaining
```



Note

The **show interface transceiver** command can only be issued on a switch in the Cisco MDS 9100 Series if the SFP is present (see [Example 12-15](#)).

Example 12-15 Displays Transceiver Information

```
switch# show interface transceiver
fc1/1 SFP is present
    name is CISCO-AGILENT
    part number is QFBR-5796L
    revision is
    serial number is A00162193
    fc-transmitter type is short wave laser
    cisco extended id is unknown (0x0)
...
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
fc1/9 SFP is present
  name is FINISAR CORP.
  part number is FTRJ-1319-7D-CSC
  revision is
  serial number is H11A6ER
  fc-transmitter type is long wave laser cost reduced
  cisco extended id is unknown (0x0)
...
```

Example 12-16 displays the running configuration for a specified interface.

Example 12-16 Displays the Running Configuration for a Specified Interface

```
switch# show running-config interface fc1/1
interface fc1/1
switchport mode FL
no shutdown
```

Displaying TL Port Information

Private loop devices refer to legacy devices that reside on arbitrated loops. These devices are not aware of a switch fabric because they only communicate with devices on the same physical loop.

The legacy devices are used in Fibre Channel networks and devices outside the loop may need to communicate with them. The communication functionality is provided through TL ports.

Use the **switchport mode** command to configure a TL port (see the “[Interface Modes](#)” section on [page 12-10](#)).

The **show tlport** command displays the TL port interface configurations. This command provides a list of all TL ports configured in a switch and shows the associated VSAN, the FC ID for the port (only domain and area are valid), and the current operational state of the TL port (up or initializing). See Examples [12-17](#) to [12-20](#).

Example 12-17 Displays the TL Ports in All VSANs

```
switch# show tlport list
-----
Interface Vsan FC-ID   State
-----
fc1/16    1      0x420000 Init
fc2/26    1      0x150000 Up
```

TL ports allow a private device (devices that physically reside on the loop) to see a fabric device and vice-versa by proxying fabric devices on the loop. Fabric devices are proxied by allocating each fabric device an ALPA on this loop.

In addition to these proxied devices, other virtual devices (local or remote domain controller addresses) are also allocated ALPAs on the loop. A switch reserves the ALPA for its own communication with private devices, and the switch acts as a SCSI initiator.

The first column in the output of the **show tlport interface** command is the ALPA identity of the device on the loop. The columns that follow include the port WWNs, the node WWNs for each device, the device as a SCSI initiator or target, and the real FC ID of the device.

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 12-18 Displays the Detailed Information for a Specific TL Port

```
switch# show tlport interface fc1/16 all
fc1/16 is up, vsan 1, FCID 0x420000
-----
alpha pWWN                                nWWN                                SCSI Type Device  FC-ID
-----
0x01 20:10:00:05:30:00:4a:de 20:00:00:05:30:00:4a:de Initiator Proxied 0xffffc42
0x73 22:00:00:20:37:39:ae:54 20:00:00:20:37:39:ae:54 Target Private 0x420073
0xef 20:10:00:05:30:00:4a:de 20:00:00:05:30:00:4a:de Initiator Switch 0x0000ef
```

Example 12-19 Displays TL Port Information for Private Devices

```
switch# show tlport int fc1/16 pri
fc1/16 is up, vsan 1, FCID 0x420000
-----
alpha pWWN                                nWWN                                SCSI Type FC-ID
-----
0x73 22:00:00:20:37:39:ae:54 20:00:00:20:37:39:ae:54 Target 0x420073
0x74 22:00:00:20:37:38:d3:de 20:00:00:20:37:38:d3:de Target 0x420074
```

Example 12-20 Displays TL Port Information for Proxied Devices

```
switch# show tlport int fc1/16 prox
fc1/16 is up, vsan 1, FCID 0x420000
-----
alpha pWWN                                nWWN                                SCSI Type FC-ID
-----
0x01 20:10:00:05:30:00:4a:de 20:00:00:05:30:00:4a:de Initiator 0xffffc42
0x02 21:00:00:e0:8b:01:95:e7 20:00:00:e0:8b:01:95:e7 Initiator 0x420100
```

TL Port Translation Guidelines

Table 12-6 lists the TL port translations supported in Cisco MDS 9000 Family switches.


Table 12-6 Supported TL Port Translations

| Translation from | Translation to | Example ¹ |
|----------------------------|-------------------------|-----------------------------|
| Private initiator | Private target | From I1 to T1 or vice versa |
| Private initiator | Public target — N port | From I1 to T2 or vice versa |
| Private initiator | Public target — NL port | From I4 to T3 or vice versa |
| Public initiator — N port | Private target | From I2 to T1 or vice versa |
| Public initiator — NL port | Private target | From I3 to T1 or vice versa |

1. See Figure 12-4.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 12-4 TL Port Translation Support Examples



91699

Follow these guidelines when configuring private loops:

- A maximum of 64 fabric devices can be proxied to a private loop.
- Fabric devices must be in the same zone as private loop devices to be proxied to the private loop.
- Each private device on a TL port may be included in a different zone.
- All devices on the loop are treated as private loops. You cannot mix private and public devices on the loop if the configured port mode is TL.
- The only FC4-type supported by TL ports is SCSI (FCP).
- Communication between a private initiator to a private target on the same private loop does not invoke TL port services.

Send documentation comments to mdsfeedback-doc@cisco.com.

Default Settings

Table 12-7 lists the default settings for Fibre Channel interface parameters.

Table 12-7 **Default Interface Parameters**

| Parameters | Default |
|------------------------------|--|
| Interface mode | Auto |
| Interface speed | Auto |
| Administrative state | Shutdown (unless changed during initial setup) |
| Trunk mode | On (unless changed during initial setup) |
| Trunk-allowed VSANs | 1 to 4093 |
| Interface VSAN | Default VSAN (1) |
| Beacon mode | Off (disabled) |
| EISL encapsulation | Disabled |
| Data field size | 2112 bytes |
| CIM server | Disabled |
| CIM server security protocol | HTTP |

Send documentation comments to mdsfeedback-doc@cisco.com.



Configuring Trunking


This chapter describes the trunking feature provided in Cisco MDS 9000 switches. It includes the following sections:

- [About Trunking, page 13-1](#)
- [About the Trunking Protocol, page 13-2](#)
- [Configuring Trunk Mode, page 13-2](#)
- [Trunk-Allowed VSAN Configuration, page 13-3](#)
- [Trunking Configuration Guidelines, page 13-6](#)
- [Displaying Trunking Information, page 13-7](#)
- [Default Settings, page 13-8](#)

About Trunking

Trunking, also known as VSAN trunking, is a feature specific to switches in the Cisco MDS 9000 Family. Trunking enables interconnect ports to transmit and receive frames in more than one VSAN, over the same physical link, using Enhanced ISL (EISL) frame format (see [Figure 13-1](#)).

Figure 13-1 Trunking



The trunking feature includes the following restrictions:

- Trunking configurations are only applicable to E ports. If trunk mode is enabled in an E port and that port becomes operational as a trunking E port, it is referred to as a TE port.
- The trunk-allowed VSANs configured for TE ports are used by the trunking protocol to determine the allowed-active VSANs in which frames can be received or transmitted.
- If a trunking enabled E port is connected to a third-party switch, the trunking protocol ensures seamless operation as an E port.

Send documentation comments to mdsfeedback-doc@cisco.com.

About the Trunking Protocol

The trunking protocol is important for E-port and TE-port operations. It supports the following:

- Dynamic negotiation of operational trunk mode.
- Selection of a common set of trunk-allowed VSANs.
- Detection of a VSAN mismatch across an ISL.

By default, the trunking protocol is enabled. If the trunking protocol is disabled on a switch, no port on that switch can apply new trunk configurations. Existing trunk configurations are not affected—the TE port continues to function in trunk mode, but only supports traffic in VSANs that it negotiated with previously (when the trunking protocol was enabled). Also, other switches that are directly connected to this switch are similarly affected on the connected interfaces. In some cases, you may need to merge traffic from different port VSANs across a non-trunking ISL. If so, disable the trunking protocol.



Tip

To avoid inconsistent configurations, shut all E ports before enabling or disabling the trunking protocol.

Enabling or Disabling the Trunking Protocol

To enable or disable the trunking protocol, follow these steps:

| | Command | Purpose |
|--------|--|--------------------------------------|
| Step 1 | switch# conf t | Enters configuration mode. |
| Step 2 | switch(config)# no trunk protocol enable switch(config)# | Disables the trunking protocol. |
| | switch(config)# trunk protocol enable switch(config)# | Enables trunking protocol (default). |

Configuring Trunk Mode

By default, the trunk mode is enabled in all Fibre Channel interfaces. However, the trunk mode configuration takes effect only in E-port mode. You can configure the trunk mode as on (enabled), off (disabled), or auto (automatic). The default trunk mode is on. The trunk mode configuration at the two ends of an ISL, between two switches, determine the resulting trunking state of the link and the port modes at both ends (see [Table 13-2](#)).

Table 13-2 *Trunk Mode Status Between Switches*

| Your Trunk Mode Configuration | | Resulting State and Port Mode | |
|-------------------------------|------------------|-------------------------------|-----------|
| Switch 1 | Switch 2 | Trunking State | Port Mode |
| On | Auto or on | Trunking (EISL) | TE port |
| Off | Auto, on, or off | No trunking (ISL) | E port |
| Auto | Auto | No trunking (ISL) | E port |

Send documentation comments to mdsfeedback-doc@cisco.com.



Tip

The preferred configuration on the Cisco MDS 9000 Family switches is one side of the trunk set to auto and the other set to on.



Note

When connected to a third-party switch, the trunk mode configuration has no effect—the ISL is always in a trunking disabled state.

Configuring the Trunk Mode

To configure the trunk mode, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/1 switch(config-if)# | Configures the specified interface. |
| Step 3 | switch(config-if)# switchport trunk mode on | Enables (default) the trunk mode for the specified interface. |
| | switch(config-if)# switchport trunk mode off | Disables the trunk mode for the specified interface. |
| | switch(config-if)# switchport trunk mode auto | Configures the trunk mode to auto mode, which provides automatic sensing for the interface. |

Trunk-Allowed VSAN Configuration


Each Fibre Channel interface has an associated trunk-allowed VSAN list. In TE-port mode, frames are transmitted and received in one or more VSANs specified in this list. By default, the VSAN range (1 through 4093) is included in the trunk-allowed list.

The common set of VSANs that are configured and active in the switch are included in the trunk-allowed VSAN list for an interface, and they are called *allowed-active* VSANs. The trunking protocol uses the list of allowed-active VSANs at the two ends of an ISL to determine the list of operational VSANs in which traffic is allowed.

In [Figure 13-1](#), switch 1 has VSANs 1 through 5, switch 2 has VSANs 1 through 3, and switch 3 has VSANs 1, 2, 4, and 5 with a default configuration of trunk-allowed VSANs. All VSANs configured in all three switches are allowed-active. However, only the common set of allowed-active VSANs at the ends of the ISL become operational as shown in [Figure 13-1](#).

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 13-1 Default Allowed-Active VSAN Configuration



You can configure a select set of VSANs (from the allowed-active list) to control access to the VSANs specified in a trunking ISL.


Using [Figure 13-1](#) as an example, you can configure the list of allowed VSANs on a per-interface basis (see [Figure 13-2](#)). For example, if VSANs 2 and 4 are removed from the allowed VSAN list of ISLs connecting to switch 1, the operational allowed list of VSANs for each ISL would be as follows:

- The ISL between switch 1 and switch 2 shall include VSAN 1 and VSAN 3.
- The ISL between switch 2 and switch 3 shall include VSAN 1 and VSAN 2.
- The ISL between switch 3 and switch 1 shall include VSAN 1, 2, and 5.

Consequently, VSAN 2 can only be routed from switch 1 through switch 3 to switch 2.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 13-2 Operational and Allowed VSAN Configuration



Configuring an Allowed-Active List of VSANs

To configure an allowed-active list of VSANs for an interface, follow these steps:


| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/1 switch(config-if)# | Configures the specified interface. |
| Step 3 | switch(config-if)# switchport trunk allowed vsan 2-4 | Changes the allowed list for the specified VSANs. |
| | switch(config-if)# switchport trunk allowed vsan add 5 updated trunking membership | Expands the specified VSAN (5) to the new allowed list. |
| | switch(config-if)# no switchport trunk allowed vsan 2-4 | Deletes VSANs 2, 3, and 4. |
| | switch(config-if)# no switchport trunk allowed vsan add 5 | Deletes the expanded allowed list. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Trunking Configuration Guidelines


If you misconfigure VSAN configurations across E ports, you could face consequences such as merging the traffic in two VSANs (thus causing both VSANs to mismatch). The trunking protocol validates the VSAN interfaces at both ends of an ISL to avoid VSANs merging (see [Figure 13-3](#)).

Figure 13-3 VSAN Mismatch



In this example, the trunking protocol detects potential VSAN merging and isolates the ports involved. The trunking protocol cannot detect merging of VSANs when a third-party switch is placed in between two Cisco MDS 9000 Family switches (see [Figure 13-4](#)).

Figure 13-4 Third-Party Switch VSAN Mismatch



VSANs 2 and 3 get effectively merged with overlapping entries in the name server and the zone applications. The Cisco MDS 9000 Fabric Manager helps detect such topologies.

Refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*.

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying Trunking Information

The **show interface** command is invoked from the EXEC mode and displays trunking configurations for a TE port. Without any arguments, this command displays the information for all of the configured interfaces in the switch. See Examples 13-1 to 13-3.

Example 13-1 Displays a Trunked Fibre Channel Interface

```
switch# show interface fc1/13
fc1/13 is trunking
  Hardware is Fibre Channel
  Port WWN is 20:0d:00:05:30:00:58:1e
  Peer port WWN is 20:0d:00:05:30:00:59:1e
  Admin port mode is auto, trunk mode is on
  Port mode is TE
  Port vsan is 1
  Speed is 2 Gbps
  Receive B2B Credit is 255
  Beacon is turned off
  Trunk vsans (admin allowed and active) (1)
  Trunk vsans (up) (1)
  Trunk vsans (isolated) ( )
  Trunk vsans (initializing) ( )
  5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    233996 frames input, 14154208 bytes, 0 discards
      0 CRC, 0 unknown class
      0 too long, 0 too short
    236 frames output, 13818044 bytes, 0 discards
    11 input OLS, 12 LRR, 10 NOS, 28 loop inits
    34 output OLS, 19 LRR, 17 NOS, 12 loop inits
```

Example 13-2 Displays the Trunking Protocol

```
switch# show trunk protocol
Trunk protocol is enabled
```

Example 13-3 Displays Per VSAN Information on Trunk Ports

```
switch# show interface trunk vsan 1-1000
fc3/1 is not trunking
...
fc3/7 is trunking
  Vsan 1000 is down (Isolation due to vsan not configured on peer)
...
fc3/10 is trunking
  Vsan 1 is up, FCID is 0x760001
  Vsan 2 is up, FCID is 0x6f0001
...
fc3/11 is trunking
  Belongs to port-channel 6
  Vsan 1 is up, FCID is 0xef0000
  Vsan 2 is up, FCID is 0xef0000
...
port-channel 6 is trunking
  Vsan 1 is up, FCID is 0xef0000
  Vsan 2 is up, FCID is 0xef0000
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Default Settings

Table 13-3 lists the default settings for trunking parameters.

Table 13-3 ***Default Trunk Configuration Parameters***

| Parameters | Default |
|------------------------|----------------------------------|
| Switch port trunk mode | On. |
| Allowed VSAN list | 1 to 4093 user-defined VSAN IDs. |
| Trunking protocol | Enabled. |



Configuring PortChannels

PortChannels refer to the aggregation of multiple physical interfaces into one logical interface to provide higher aggregated bandwidth, load balancing, and link redundancy. PortChannels can connect to interfaces across switching modules, so a failure of a switching module cannot bring down the PortChannel link.

This chapter discusses the PortChannel feature provided in the switch and includes the following sections:

- [PortChannel Functionality, page 14-2](#)
- [PortChannel Examples, page 14-2](#)
- [About PortChanneling and Trunking, page 14-4](#)
- [About Load Balancing, page 14-4](#)
- [PortChannel Creation, page 14-6](#)
- [PortChannel Modes, page 14-7](#)
- [Deleting PortChannels, page 14-8](#)
- [Interface Addition to a PortChannel, page 14-8](#)
- [Deleting Interfaces from a PortChannel, page 14-11](#)
- [PortChannel Configuration Guidelines, page 14-11](#)
- [PortChannel Protocol, page 14-13](#)
- [PortChannel Configuration Verification, page 14-17](#)
- [Default Settings, page 14-20](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

PortChannel Functionality

A PortChannel has the following functionality:

- Provides a point-to-point connection over ISL (E ports) or EISL (TE ports). Multiple links can be combined into a PortChannel.
- Increases the aggregate bandwidth on an ISL by distributing traffic among all functional links in the channel.
- Load balances across multiple links and maintains optimum bandwidth utilization. Load balancing is based on the source ID, destination ID, and exchange ID (OX ID).
- Provides high availability on an ISL. If one link fails, traffic previously carried on this link is switched to the remaining links. If a link goes down in a PortChannel, the upper protocol is not aware of it. To the upper protocol, the link is still there, although the bandwidth is diminished. The routing tables are not affected by link failure. PortChannels may contain up to 16 physical links and may span multiple modules for added high availability.

**Note**

See the [“Fail-Over Scenarios for PortChannels and FSPF Links”](#) section on page 24-3 for fail-over scenarios.

Cisco MDS 9000 Family of switches support 128 PortChannels with 16 interfaces per PortChannel. A PortChannel number refers to the unique (to each switch) identifier associated with each channel group. This number ranges from 1 to 128.


PortChannel Examples

PortChannels on Cisco MDS 9000 Family switches allow flexibility in configuration. [Figure 14-1](#) illustrates three possible PortChannel configurations:

- PortChannel A aggregates two links on two interfaces on the same switching module at each end of a connection.
- PortChannel B also aggregates two links, but each link is connected to a different switching module. If the switching module goes down, traffic is not affected.
- PortChannel C aggregates three links. Two links are on the same switching module at each end, while one is connected to a different switching module on switch 2.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 14-1 PortChannel Flexibility



32-Port Switching Module Configuration Guidelines

The 32-port switching module guidelines applies to the following hardware:

- The 32-port 2-Gbps or 1-Gbps switching module
- The Cisco MDS 9140 Switch

When configuring these host-optimized ports, the following PortChannel guidelines apply:

- Any (or all) full line rate port(s) in the Cisco MDS 9100 Series can be included in a PortChannel.
- The host-optimized ports in the Cisco MDS 9100 Series are subject to the same PortChannel rules as 32-port switching modules—only the first port of each group of 4 ports is included in a PortChannel.
 - You can configure only the first port in each 4-port group (for example, the first port in ports 1-4, the fifth port in ports 5-8 and so on) as an E port. If the first port in the group is configured as a PortChannel, the other three ports in each group (ports 2-4, 6-8 and so on) are not usable and remain in the shutdown state.
 - If any of the other three ports are configured in a no shutdown state, you cannot configure the first port to be a PortChannel. The other three ports continue to remain in a no shutdown state.



Note

In the Cisco MDS 9100 Series, the left most groups of ports outlined in white (4 ports in the Cisco MDS 9120 Switch and 8 ports in the Cisco MDS 9140 Switch) are full line rate like the 16-port switching module. The other ports (16 ports in the Cisco MDS 9120 Switch and 32 ports in the Cisco MDS 9140 Switch) are host-optimized like the 32-port switching module. Each group of 4 host-optimized ports have the same rules as for the 32-port switching module.

Send documentation comments to mdsfeedback-doc@cisco.com.

About PortChanneling and Trunking

Trunking is a commonly-used storage industry term. However, the Cisco SAN-OS software and switches in the Cisco MDS 9000 Family implement trunking and PortChanneling as defined below:

- PortChanneling enables several physical links to be combined into one aggregated logical link.
- Trunking enables a link transmitting frames in the EISL format to carry (trunk) multiple VSAN traffic. When trunking is operational on an E port, that E port becomes a TE port. A TE port is specific to switches in the Cisco MDS 9000 Family. An industry standard E port can link to other vendor switches and is referred to as a nontrunking interface (see [Figure 14-2](#) and [Figure 14-3](#)).

See [Chapter 13, “Configuring Trunking”](#) for information on trunked interfaces.

Figure 14-2 Trunking Only





Figure 14-3 PortChanneling and Trunking



PortChanneling and trunking are used separately across an ISL:

- PortChanneling—Interfaces can be channeled between E ports and TE ports.
- Trunking—Trunking, which permits carrying traffic on multiple VSANs between switches, can be done only between TE ports.

See [Chapter 10, “Configuring and Managing VSANs.”](#)

Both PortChanneling and trunking can be used between TE ports over EISLs.

About Load Balancing

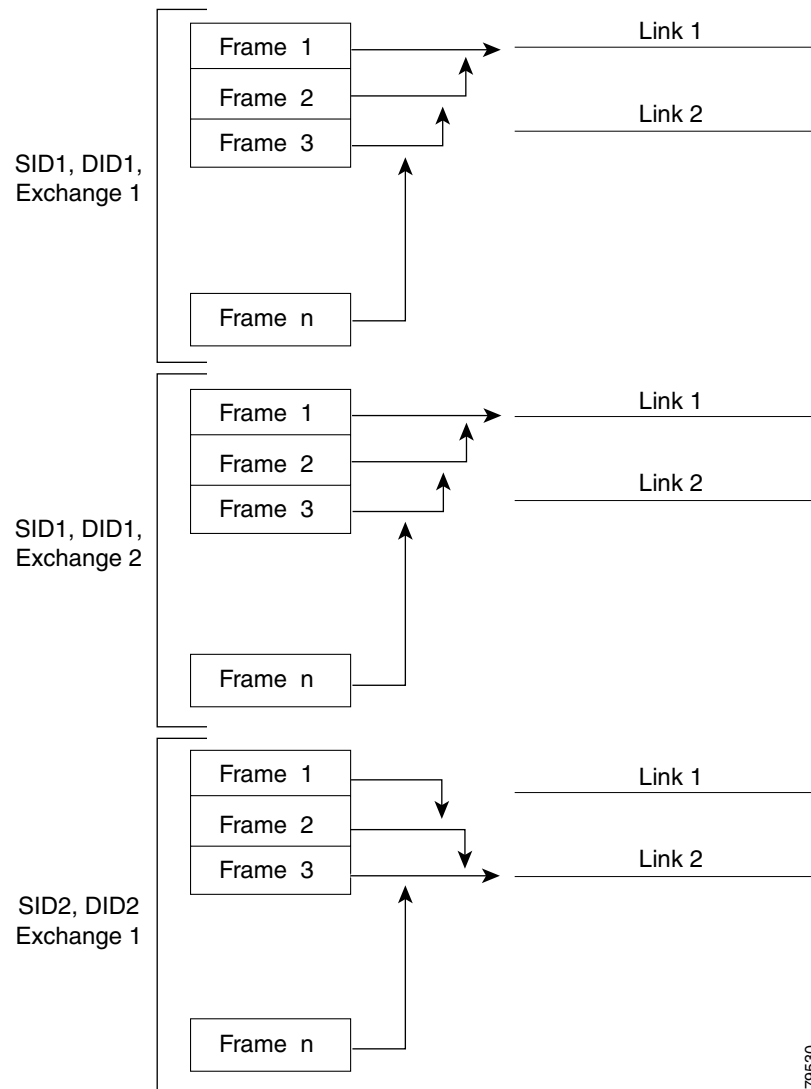
Two mechanisms support the load balancing functionality:

- Flow based—All frames between source and destination follow the same links for a given flow. That is, whichever link is selected for the first exchange of the flow is used for all subsequent exchanges.
- Exchange based—The first frame in an exchange picks a link and subsequent frames in the exchange follow the same link. However, subsequent exchanges can use a different link. This provides more granular load balancing while preserving the order of frames for each exchange.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 14-4 illustrates how source ID 1 (SID1) and destination ID1-based(DID1) load balancing works. When the first frame in a flow is received on an interface for forwarding, link 1 is selected. Each subsequent frame in that flow is sent over the same link. No frame in SID1 and DID1 utilizes link 2.

Figure 14-4 SID1 and DID1Based Load Balancing

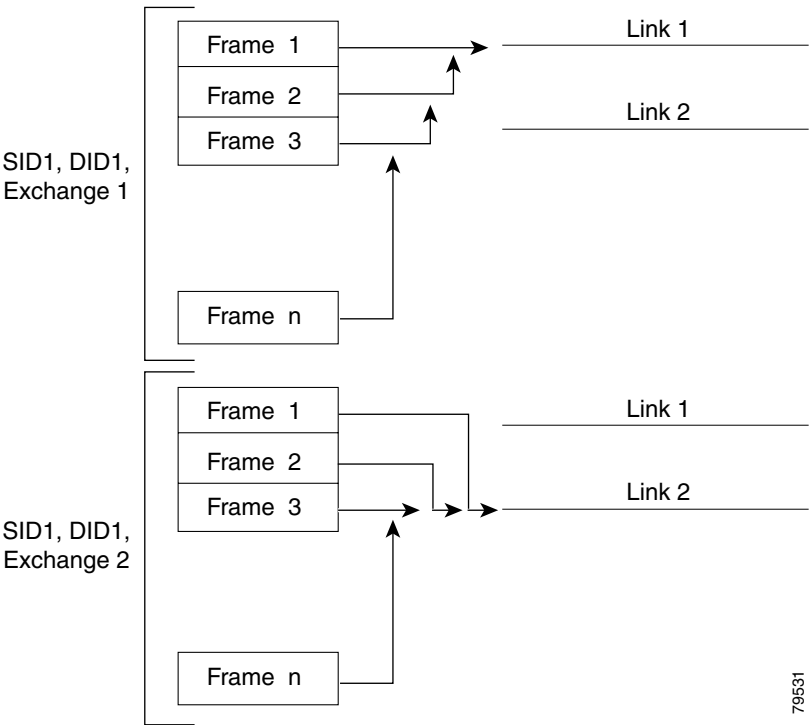


79530

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 14-5 illustrates how exchange based load balancing works. When the first frame in an exchange is received for forwarding on an interface, link 1 is chosen by a hash algorithm. All remaining frames in that particular exchange are sent on the same link. For exchange 1, no frame uses link 2. For the next exchange, link 2 is chosen by the hash algorithm. Now all frames in exchange 2 use link 2.

Figure 14-5 SID1, DID1, and Exchange Based Load Balancing



For more information on configuring load balancing and in-order delivery features, see the [“VSAN Attributes” section on page 10-6](#).

PortChannel Creation

PortChannels are created with default values. You can change the default configuration just like any other physical interface.

To create a PortChannel, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface port-channel 1 switch(config-if)# | Configures the specified PortChannel (1) using the default ON mode. |

Send documentation comments to mdsfeedback-doc@cisco.com.

PortChannel Modes

You can configure each PortChannel with a channel group mode parameter to determine the PortChannel protocol behavior for all member ports in this channel group. The possible values for a channel group mode are as follows.

- **ON (default)**—The member ports only operate as part of a PortChannel or remain inactive. In this mode, the PortChannel protocol is not initiated. However, if a PortChannel protocol frame is received from a peer port, the software indicates its nonnegotiable status. This mode is backward compatible with the existing implementation of PortChannels in releases prior to Release 2.0(1b), where the channel group mode is implicitly assumed to be ON. In Cisco MDS SAN-OS Releases 1.3 and earlier, the only available PortChannel mode was the ON mode. PortChannels configured in the ON mode require you to explicitly enable and disable the Portchannel member ports at either end if you add or remove ports from the PortChannel configuration. You must physically verify that the local and remote ports are connected to each other.
- **ACTIVE**—The member ports initiate PortChannel protocol negotiation with the peer port(s) regardless of the channel group mode of the peer port. If the peer port, while configured in a channel group, does not support the PortChannel protocol, or responds with a nonnegotiable status, it will default to the ON mode behavior. As of Cisco MDS SAN-OS Release 2.0(1b), the ACTIVE PortChannel mode is introduced to enable an automatic recovery without explicitly enabling and disabling the PortChannel member ports at either end.

Table 14-2 a comparative reference for both modes.

Table 14-1 Channel Group Configuration Differences

| ON Mode | ACTIVE Mode |
|---|---|
| No protocol is exchanged. | A PortChannel protocol negotiation is performed with the peer ports. |
| Moves interfaces to the suspended state if its operational values are incompatible with the PortChannel. | Moves interfaces to the isolated state if its operation values are incompatible with the PortChannel. |
| When you add or modify a PortChannel member port configuration, you must explicitly disable (shut) and enable (no shut) the PortChannel member ports at either end. | When you add or modify a PortChannel interface, the PortChannel automatically recovers. |
| Port initialization is not synchronized. | Synchronized bringup of all ports in a channel across peer switches. |
| All misconfigurations are not detected as no protocol is exchanged. | Consistently detect misconfigurations using a PortChannel protocol. |
| Transitions misconfigured ports to suspended state, you must explicitly disable (shut) and enable (no shut) the member ports at either end. | Transitions misconfigured ports to isolated state to correct the misconfigurations. Once you correct the misconfiguration, the protocol ensures automatic recovery. |
| This is the default mode | You must explicitly configure this mode. |

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure the active mode, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface port-channel 1 switch(config-if)# | Configures the specified PortChannel (1) using the default ON mode. |
| Step 3 | switch(config-if)# channel mode active | Configures the ACTIVE mode. |
| | switch(config-if)# no channel mode active | Reverts to the default ON mode. |

Deleting PortChannels

When you delete the PortChannel, the corresponding channel membership is also deleted. All interfaces in the deleted PortChannel convert to individual physical links. After the PortChannel is removed, regardless of the mode (ACTIVE and ON) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the [“Graceful Shut Down” section on page 12-9](#)).

If you delete the PortChannel for one port, then the individual ports within the deleted PortChannel retain the compatibility parameter settings (speed, mode, port VSAN, allowed VSAN, and port security). You can explicitly change those settings as required.

- If you use the default ON mode, to avoid inconsistent states across switches, and to maintain consistency across switches, the ports shut down. You must explicitly enable those ports again.
- If you use the ACTIVE mode, the PortChannel ports automatically recover from the deletion.

To delete a PortChannel, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# no interface port-channel 1 port-channel 1 deleted and all its members disabled please do the same operation on the switch at the other end of the port-channel switch(config)# | Deletes the specified PortChannel (1), its associated interface mappings, and the hardware associations for this PortChannel. |

Interface Addition to a PortChannel

You can add a physical interface (or a range of interfaces) to a nonexistent or an existing PortChannel and the PortChannel is automatically created. If the PortChannel does not exist, it is created. The compatible parameters on the configuration are mapped to the PortChannel. Adding an interface to a PortChannel increases the channel size and bandwidth of the PortChannel.

After the members are added, regardless of the mode (ACTIVE and ON) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the [“Graceful Shut Down” section on page 12-9](#)).

Send documentation comments to mdsfeedback-doc@cisco.com.

To add a port to a PortChannel, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/15 switch(config-if)# | Configures the specified port interface (fc1/15). |
| Step 3 | switch(config-if)# channel-group 15 fc1/15 added to port-channel 15 and disabled please do the same operation on the switch at the other end of the port-channel, then do "no shutdown" at both ends to bring them up | Adds physical Fibre Channel port 1/15 to channel group 15. If channel group 15 does not exist, it is created. The port is shut down. |

To add a range of ports to a PortChannel, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/1 - 5 switch(config-if)# | Configures the specified range of interfaces. In this example, interfaces from 1/1 to 1/5 are configured. |
| Step 3 | switch(config-if)# channel-group 2 fc1/1 fc1/2 fc1/3 fc1/4 fc1/5 added to port-channel 2 and disabled please do the same operation on the switch at the other end of the port-channel, then do "no shutdown" at both ends to bring them up | Adds physical interfaces 1/1, 1/2, 1/3, 1/4, and 1/5 to channel group 2. If channel group 2 does not exist, it is created. If the compatibility check is successful, the interfaces are operational and the corresponding states apply to these interfaces. |

Forcing an Interface Addition

You can specify a **force** option to force the port configuration to be overwritten by the PortChannel. In this case, the interface is added to a PortChannel.

- If you use the default ON mode, to avoid inconsistent states across switches, and to maintain consistency across switches, the ports shut down. You must explicitly enable those ports again.
- If you use the ACTIVE mode, the PortChannel ports automatically recover from the addition.



Note

When PortChannels are created from within an interface, the **force** option cannot be used.

After the members are forcefully added, regardless of the mode (ACTIVE and ON) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the [“Graceful Shut Down”](#) section on page 12-9).

Send documentation comments to mdsfeedback-doc@cisco.com.

To force the addition of a port to a PortChannel, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/1 switch(config-if)# | Configures the specified port interface (fc1/1). |
| Step 3 | switch(config-if)# channel-group 1 force fc1/1 added to port-channel 1 and disabled please do the same operation on the switch at the other end of the port-channel, then do "no shutdown" at both ends to bring them up switch(config-if)# | Forces a physical Fibre Channel port 1/1 addition to channel group 1. The E port is shut down. |

Compatibility Check

A compatibility check ensures that the same parameter settings are used in all physical ports in the channel. Otherwise, they cannot become part of a PortChannel. The compatibility check is performed before a port is added to the PortChannel.

The check ensures that the following parameters and setting match at both ends of a PortChannel:

- Capability parameters (type of interface, Gigabit Ethernet at both ends or Fibre Channel at both ends).
- Administrative compatibility parameters (speed, mode, port VSAN, allowed VSAN, and port security).
- Operational parameters (speed and remote switch's WWN).

A port addition procedure fails if the capability and administrative parameters in the remote switch are incompatible with the capability and administrative parameters in the local switch. If the compatibility check is successful, the interfaces are operational and the corresponding compatibility parameter settings apply to these interfaces.

Suspended and Isolated States

If the operational parameters are incompatible, the compatibility check fails and the interface is placed in a suspended or isolated state based on the configured mode:

- An interface enters the suspended state if the interface is configured in the ON mode.
- An interface enters the isolated state if the interface is configured in the ACTIVE mode.

See the ["Reason Codes" section on page 12-6](#).

Send documentation comments to mdsfeedback-doc@cisco.com.

Deleting Interfaces from a PortChannel

When a physical interface is deleted from the PortChannel, the channel membership is automatically updated. If the deleted interface is the last operational interface, then the PortChannel status is changed to a down state. Deleting an interface from a PortChannel decreases the channel size and bandwidth of the PortChannel.

- If you use the default ON mode, to avoid inconsistent states across switches, and to maintain consistency across switches, the ports shut down. You must explicitly enable those ports again.
- If you use the ACTIVE mode, the PortChannel ports automatically recover from the deletion.

After the members are deleted, regardless of the mode (ACTIVE and ON) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the [“Graceful Shut Down”](#) section on page 12-9).

To delete a physical interface (or a range of physical interfaces), follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch(config)# interface fc1/1 switch(config-if)# | Enters the selected physical interface level. |
| | switch(config)# interface fc1/1 - 5 switch(config-if)# | Enters the selected range of physical interfaces. |
| Step 2 | switch(config-if)# no channel-group 2 fc1/1 fc1/2 fc1/3 fc1/4 fc1/5 removed from port-channel 2 and disabled. Please do the same operation on the switch at the other end of the port-channel switch(config-if)# | Deletes the physical Fibre Channel interfaces in channel group 2. |

PortChannel Configuration Guidelines

Before configuring a PortChannel, consider the following guidelines

- Configure the PortChannel across switching modules to prevent redundancy on switching module reboots or upgrades.
- Ensure that one PortChannel is not connected to different sets of switches. PortChannels require point-to-point connections between the same set of switches.

Error Detection

If you misconfigure PortChannels, you may receive a misconfiguration message. If you receive this message, the PortChannel's physical links are disabled since an error has been detected.

A PortChannel error is detected if the following requirements are not met:

- Each switch on either side of a PortChannel must be connected to the same number of interfaces.
- Each interface must be connected to a corresponding interface on the other side (see the [“Invalid Configuration Examples”](#) section on page 14-13).
- Links in a PortChannel cannot be changed after the PortChannel is configured. If you change the links after the PortChannel is configured, be sure to reconnect the links to interfaces within the PortChannel and re-enable the links.

If all three conditions are not met, the faulty link is disabled.

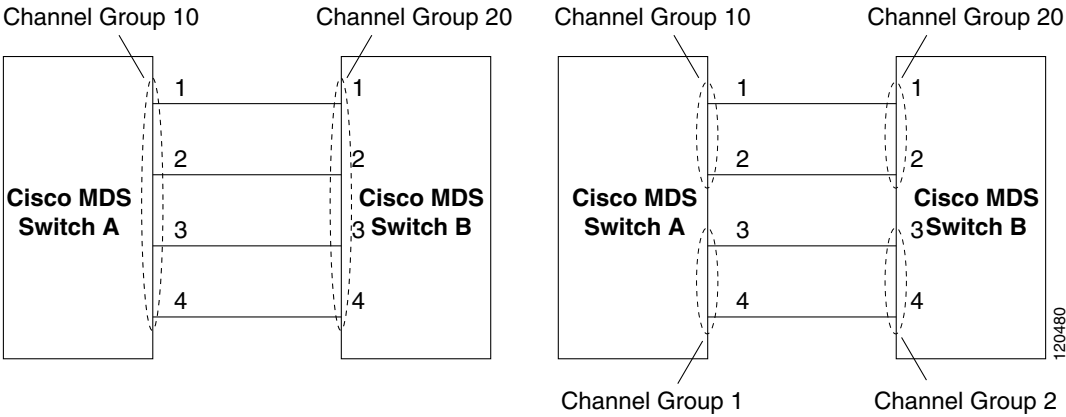
Send documentation comments to mdsfeedback-doc@cisco.com.

Issue the **show interface** command for that interface to verify that the PortChannel is functioning as required.

Valid Configurations

Figure 14-6 provides examples of valid configurations.

Figure 14-6 Valid Configurations

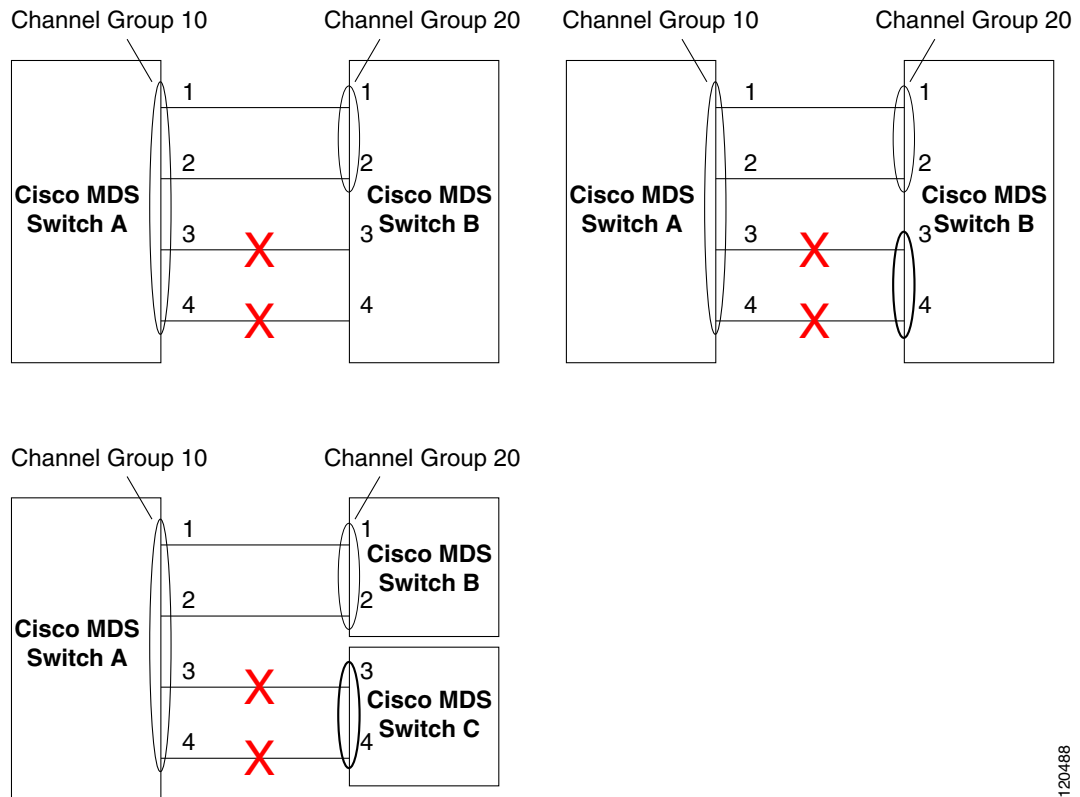


Send documentation comments to mdsfeedback-doc@cisco.com.

Invalid Configuration Examples

Figure 14-7 provides examples of invalid configurations. Assuming that the links are brought up in the 1,2,3, 4 sequence, links 3 and 4 will be operationally down as the fabric is misconfigured.

Figure 14-7 Misconfigured Configurations



120488

PortChannel Protocol

In earlier Cisco SAN-OS releases, PortChannels required additional administrative tasks to support synchronization. The Cisco SAN-OS software now provides more robust error detection and synchronization capabilities. You can manually configure channel groups or they can be automatically created. In both cases, the channel group have the same capability and configurational parameters. Any change in configuration applied to the associated PortChannel interface is propagated to all members of the channel group.

As of Cisco SAN-OS Release 2.0(1b), a protocol to exchange PortChannel configurations is available in all Cisco MDS switches. This addition simplifies PortChannel management with incompatible ISLs. An additional autcreation mode enables ISLs with compatible parameters to automatically form channel groups without manual intervention.

The PortChannel protocol is enabled by default in Cisco SAN-OS Release 2.0(1b) and later.

Send documentation comments to mdsfeedback-doc@cisco.com.

About PortChannel Protocols

The PortChannel protocol expands the PortChannel functional model in Cisco MDS switches. It uses the exchange peer parameters (EPP) services to communicate across peer ports in an ISL. Each switch uses the information received from the peer ports along with its local configuration and operational values to decide if it should be part of a PortChannel. The protocol ensures that a set of ports are eligible to be part of the same PortChannel. They are only eligible to be part of the same port channel if all the ports have a compatible partner.


The PortChannel protocol uses two sub-protocols:

- Bringup protocol—automatically detects misconfigurations so you can correct them. This protocol synchronizes the PortChannel at both ends so that all frames for a given flow (as identified by the source FCID, destination FCID and OX_ID) are carried over the same physical link in both directions. This helps make applications like write acceleration work for PortChannels over FCIP links.
- Autocreation protocol—automatically aggregates compatible ports into a PortChannel.

Channel Group Creation

Assuming link A1-B1 comes up first in [Figure 14-8](#), that link is operational as an individual link. When the next link, say A2-B2 comes up, the PortChannel protocol identifies if this link is compatible with link A1-B1 and automatically creates channel groups 10 and 20 in the respective switches. If link A3-B3 can join the channel groups (and hence, the PortChannels), the respective ports have compatible configurations. If link A4-B4 operates as an individual link, it is due to the incompatible configuration of the two end ports with the other member ports in this channel group.

Figure 14-8 Auto-Creating Channel Groups



The channel group numbers are selected dynamically, and as such, the administrative configuration of the ports forming the channel group at either end are applicable to the newly created channel group. The channel group number being chosen dynamically may be different across reboots for the same set of PortChannels based on the order of ports that are initialized in the switch.

[Table 14-2](#) identifies the differences between user-configured and auto-configured channel groups.

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 14-2 Channel Group Configuration Differences

| User Configured Channel Group | Autocreated Channel Group |
|--|--|
| Manually configured by the user. | Created automatically when compatible links come up between two compatible switches, if channel group autocreation is enabled in all ports at both ends. |
| Member ports can not participate in autocreation of channel groups. The autocreation feature cannot be configured. | None of these ports are members of a user configured channel group. |
| You can form the PortChannel with a subset of the ports in the channel group. Incompatible ports remain in a suspended or isolated state depending on the ON or ACTIVE mode configuration. | All ports included in the channel group participate in the PortChannel—no member port becomes isolated or suspended; instead, the member port is removed from the channel group when the link is found to be incompatible. |
| Any administrative configuration made to the PortChannel is applied to all ports in the channel group, and you can save the configuration for the PortChannel interface. | Any administrative configuration made to the PortChannel is applied to all ports in the channel group, but the configurations are saved for the member ports; no configuration is saved for the PortChannel interface. You can explicitly convert this channel group, if required. |
| You can remove any channel group and add members to a channel group. | You cannot remove a channel group, or add/remove any of its members. The channel group is removed when no member ports exist. |

Autocreation Functionality

The autocreation protocol has the following functionality:

- A port is not allowed to be configured as part of a PortChannel when the autocreation feature is enabled. These two configurations are mutually exclusive.
- Autocreation must be enabled in both the local and peer ports to negotiate a PortChannel.
- Aggregation occurs in one of two ways:
 - A port is aggregated into a compatible autocreated PortChannel, or
 - A port is aggregated with another compatible port to form a new PortChannel
- Newly created PortChannels are allocated from the maximum possible PortChannel (128) in a decreasing order based on availability. If all 128 numbers are used up, aggregation is not allowed.
- You cannot change the membership or delete an autocreated PortChannel
- When you disable autocreation, all member ports are removed from the autocreated PortChannel.
- Once the last member is removed from an autocreated PortChannel, the channel is automatically deleted and the number is released for reuse.
- An autocreated PortChannel is not persistent through a reboot. An autocreated PortChannel can be manually configured to appear the same as a persistent PortChannel. Once the PortChannel is made persistent, the autocreation feature is disabled in all member ports.

Send documentation comments to mdsfeedback-doc@cisco.com.

Enabling and Configuring Autocreation

You can enable or disable the autocreation feature on a per-port basis or for all ports in the switch. When this configuration is enabled, the channel group mode is assumed to be active. The default for this task is disabled.

If autocreation of channel groups is enabled for an interface, you must first disable autocreation before downgrading to earlier software versions or before configuring the interface in a manually configured channel group.



Tip

When enabling autocreation in any switch in the Cisco MDS 9000 Family, we recommend that you retain at least one interconnected port between the switches without any autocreation configuration. If all ports between two switches are configured with the autocreation feature at the same time, you may face a possible traffic disruption between these two switches as the ports are automatically disabled and reenabled when ports are added to a autocreated PortChannel.

To configure automatic channel groups, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc8/13 switch(config-if)# | Enters the configuration mode for the selected interface(s). |
| Step 3 | switch(config-if)# channel-group auto | Automatically creates the channel group for the selected interface(s). |
| | switch(config-if)# no channel-group auto | Disables the autocreation of channel groups for this interface, even if the system default configuration may have autocreation enabled. |

Converting to Manually-Configured Channel Groups

A user-configured channel group cannot be converted to an autocreated channel group. However, you can convert an autocreated channel group to a manual channel group. Once performed, this task is irreversible—the channel group number does not change, but the member ports operate according to the properties of the manually configured channel group, and the autocreation of channel group is implicitly disabled for all member ports.



Tip

If you enable persistence, be sure to enable it at both ends of the PortChannel.

You can convert autocreated channel group to a user-configured channel group using the **port-channel channel-group-number persistent EXEC** command. If the PortChannel does not exist, this command is not executed.

Send documentation comments to mdsfeedback-doc@cisco.com.

PortChannel Configuration Verification

You can view specific information about existing PortChannels at any time from EXEC mode. The following **show** commands provide further details on existing PortChannels. You can force all screen output to go to a printer or save it to a file. See Examples 14-1 to 14-6.

The **show port-channel summary** command displays a summary of PortChannels within the switch. A one-line summary of each PortChannel provides the administrative state, the operational state, the number of attached and active interfaces (up), and the first operational port (FOP), which is the primary operational interface selected in the PortChannel to carry control-plane traffic (no load-balancing). The FOP is the first port that comes up in a PortChannel and can change if the port goes down. The FOP is also identified by an asterisk (*).

Example 14-1 Displays the PortChannel Summary

```
switch# show port-channel summary
```

| Interface | Total Ports | Oper Ports | First Oper Port |
|-----------------|-------------|------------|-----------------|
| port-channel 77 | 2 | 0 | -- |
| port-channel 78 | 2 | 0 | -- |
| port-channel 79 | 2 | 2 | fcip200 |

Example 14-2 Displays the PortChannel Configured in the Default ON Mode

```
switch# show port-channel database
port-channel 77
  Administrative channel mode is on
  Operational channel mode is on
  Last membership update succeeded
  2 ports in total, 0 ports up
  Ports:  fcip1    [down]
          fcip2    [down]
port-channel 78
  Administrative channel mode is on
  Operational channel mode is on
  Last membership update succeeded
  2 ports in total, 0 ports up
  Ports:  fc2/1    [down]
          fc2/5    [down]
port-channel 79
  Administrative channel mode is on
  Operational channel mode is on
  Last membership update succeeded
  First operational port is fcip200
  2 ports in total, 2 ports up
  Ports:  fcip101  [up]
          fcip200  [up] *
```

Example 14-3 Displays the PortChannel Configured in the ACTIVE Mode

```
switch# show port-channel database
port-channel 77
  Administrative channel mode is active
  Operational channel mode is active
  Last membership update succeeded
  2 ports in total, 0 ports up
  Ports:  fcip1    [down]
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```

        fcip2      [down]
port-channel 78
  Administrative channel mode is active
  Operational channel mode is active
  Last membership update succeeded
  2 ports in total, 0 ports up
  Ports:   fc2/1      [down]
           fc2/5      [down]
port-channel 79
  Administrative channel mode is active
  Operational channel mode is active
  Last membership update succeeded
  First operational port is fcip200
  2 ports in total, 2 ports up
  Ports:   fcip101    [up]
           fcip200    [up] *
```

The **show port-channel consistency** command has two options—without and with details.

Example 14-4 Displays the Consistency Status without Details

```

switch# show port-channel consistency
Database is consistent
```

Example 14-5 Displays the Consistency Status with Details

```

switch# show port-channel consistency detail
Authoritative port-channel database:
=====
totally 3 port-channels
port-channel 77:
  2 ports, first operational port is none
  fcip1      [down]
  fcip2      [down]
port-channel 78:
  2 ports, first operational port is none
  fc2/1      [down]
  fc2/5      [down]
port-channel 79:
  2 ports, first operational port is fcip200
  fcip101    [up]
  fcip200    [up]
=====
database 1: from module 5
=====
totally 3 port-channels
port-channel 77:
  2 ports, first operational port is none
  fcip1      [down]
  fcip2      [down]
port-channel 78:
  2 ports, first operational port is none
  fc2/1      [down]
  fc2/5      [down]
port-channel 79:
  2 ports, first operational port is fcip200
  fcip101    [up]
  fcip200    [up]
=====
database 2: from module 4
=====
totally 3 port-channels
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
port-channel 77:
  2 ports, first operational port is none
  fcip1    [down]
  fcip2    [down]
port-channel 78:
  2 ports, first operational port is none
  fc2/1    [down]
  fc2/5    [down]
port-channel 79:
  2 ports, first operational port is fcip200
  fcip101  [up]
  fcip200  [up]
...
```

The **show port-channel usage** command displays details of the used and unused PortChannel numbers.

Example 14-6 Displays the PortChannel Usage

```
switch# show port-channel usage
Totally 3 port-channel numbers used
=====
Used   :    77 - 79
Unused:    1 - 76 , 80 - 128
```

Example 14-7 Displays the PortChannel Compatibility

```
switch# show port-channel compatibility-parameters
physical port layer          fibre channel or ethernet
port mode                    E/AUTO only
trunk mode
speed
port VSAN
port allowed VSAN list
```

Use the existing **show** commands to obtain further details on autocreated channel group attributes. Autocreated PortChannels are indicated explicitly to help differentiate them from the manually-created PortChannels. See Examples 14-8 to 14-10.

Example 14-8 Displays Autocreated PortChannels

```
switch# show interface fc1/1
fc1/1 is trunking
Hardware is Fibre Channel, FCOT is short wave laser
Port WWN is 20:0a:00:0b:5f:3b:fe:80
...
Receive data field Size is 2112
Beacon is turned off
Port-channel auto creation is enabled
Belongs to port-channel 123
...
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 14-9 Displays the Specified PortChannel Interface

```
switch# show port-channel database interface port-channel 128
port-channel 128
  Administrative channel mode is active
  Operational channel mode is active
  Last membership update succeeded
  Channel is auto created
  First operational port is fc1/1
  1 ports in total, 1 ports up
  Ports:   fc1/1   [up] *
```

Example 14-10 Displays the PortChannel Summary

```
switch# show port-channel summary
-----
Interface                Total Ports      Oper Ports      First Oper Port
-----
port-channel 1            1                0               --
port-channel 2            1                1              fc8/13
port-channel 3            0                0               --
port-channel 4            0                0               --
port-channel 5            1                1              fc8/3
port-channel 6            0                0               --
```

Default Settings

Table 14-3 lists the default settings for PortChannels.

Table 14-3 Default PortChannel Parameters

| Parameters | Default |
|--------------------------|-----------------------------|
| PortChannels | FSPF is enabled by default. |
| Create PortChannel | Administratively up. |
| Default PortChannel mode | ON. |
| Autocreation | Disabled. |



Configuring and Managing Zones

Zoning enables you to set up access control between storage devices or user groups. If you have administrator privileges in your fabric, you can create zones to increase network security and to prevent data loss or corruption. Zoning is enforced by examining the source-destination ID field.

As of Cisco SAN-OS Release 2.0(1b), advanced zoning capabilities specified in the FC-GS-4 and FC-SW-3 standards are provided. You can use either the existing basic zoning capabilities or the advanced, standards-compliant zoning capabilities.

This chapter includes the following sections:

- [Zoning Features, page 15-2](#)
- [Zoning Example, page 15-3](#)
- [Zone Implementation, page 15-4](#)
- [Zone Configuration, page 15-4](#)
- [Zone Set Creation, page 15-7](#)
- [Zone Enforcement, page 15-10](#)
- [The Default Zone, page 15-11](#)
- [Recovering from Link Isolation, page 15-13](#)
- [Zone Set Distribution, page 15-11](#)
- [Zone Set Duplication, page 15-14](#)
- [Zone Database Information, page 15-15](#)
- [Zone-Based Traffic Priority, page 15-15](#)
- [Configuring Broadcast Zoning, page 15-17](#)
- [About LUN Zoning, page 15-17](#)
- [About Read-Only Zones, page 15-19](#)
- [Renaming Zones, Zone Sets, fcaliases, and Zone Attribute Groups, page 15-21](#)
- [Cloning Zones, Zone Sets, fcaliases, and Zone Attribute Groups, page 15-21](#)
- [Displaying Zone Information, page 15-21](#)
- [About Enhanced Zoning, page 15-27](#)
- [Displaying Enhanced Zone Information, page 15-33](#)
- [Default Settings, page 15-36](#)

Send documentation comments to mdsfeedback-doc@cisco.com.



Note

Table 10-1 on page 10-4 lists the differences between zones and VSANs.

Zoning Features

Zoning has the following features:


- A zone consists of multiple zone members.
 - Members in a zone can access each other; members in different zones cannot access each other.
 - If zoning is not activated, all devices are members of the default zone.
 - If zoning is activated, any device that is not in an active zone (a zone that is part of an active zone set) is a member of the default zone.
 - Zones can vary in size.
 - Devices can belong to more than one zone.
- A zone set consists of one or more zones.
 - A zone set can be activated or deactivated as a single entity across all switches in the fabric.
 - Only one zone set can be activated at any time.
 - A zone can be a member of more than one zone set.
- Zoning can be administered from any switch in the fabric.
 - When you activate a zone (from any switch), all switches in the fabric receive the active zone set. Additionally, full zone sets are distributed to all switches in the fabric, if this feature is enabled in the source switch.
 - If a new switch is added to an existing fabric, zone sets are acquired by the new switch.
- Zone changes can be configured nondisruptively. New zones and zone sets can be activated without interrupting traffic on unaffected ports or devices.
- Zone membership criteria is based on WWNs or FC IDs.
 - Port world wide name (pWWN)—Specifies the pWWN of an N port attached to the switch as a member of the zone.
 - Fabric pWWN—Specifies the WWN of the fabric port (switch port's WWN). This membership is also referred to as port-based zoning.
 - FC ID—Specifies the FC ID of an N port attached to the switch as a member of the zone.
 - Interface and switch WWN (sWWN)—Specifies the interface of a switch identified by the sWWN. This membership is also referred to as interface-based zoning.
 - Interface and domain ID—Specifies the interface of a switch identified by the domain ID.
 - Domain ID and port number—Specifies the domain ID of an MDS domain and additionally specifies a port belonging to a non-Cisco switch.
 - IP address—Specifies the IP address (and optionally the subnet mask) of an attached device.
- Default zone membership includes all ports or WWNs that do not have a specific membership association. Access between default zone members is controlled by the default zone policy.

Send documentation comments to mdsfeedback-doc@cisco.com.

Zoning Example


Figure 15-1 illustrates a zone set with two zones, zone 1 and zone 2, in a fabric. Zone 1 provides access from all three hosts (H1, H2, H3) to the data residing on storage systems S1 and S2. Zone 2 restricts the data on S3 to access only by H3. Note that H3 resides in both zones.

Figure 15-1 ***Fabric with Two Zones***



Of course, there are other ways to partition this fabric into zones. Figure 15-2 illustrates another possibility. Assume that there is a need to isolate storage system S2 for the purpose of testing new software. To achieve this, zone 3 is configured, which contains only host H2 and storage S2. You can restrict access to just H2 and S2 in zone 3, and to H1 and S1 in zone 1.

Figure 15-2 ***Fabric with Three Zones***



Send documentation comments to mdsfeedback-doc@cisco.com.

Zone Implementation

All switches in the Cisco MDS 9000 Family automatically support the following basic zone features (no additional configuration is required):

- Zones are contained in a VSAN.
- Hard zoning cannot be disabled.
- Name server queries are soft-zoned.
- Only active zone sets are distributed.
- Unzoned devices cannot access each other.
- A zone or zone set with the same name can exist in each VSAN.
- Each VSAN has a full database and an active database.
- Active zone sets cannot be changed, without activating a full zone database.
- Active zone sets are preserved across switch reboots.
- Changes to the full database must be explicitly saved.
- Zone reactivation (a zone set is active and you activate another zone set) does not disrupt existing traffic.

If required, you can additionally configure the following zone features:

- Propagate full zone sets to all switches on a per VSAN basis.
- Change the default policy for unzoned members.
- Interoperate with other vendors by configuring a VSAN in the interop mode. You can also configure one VSAN in the interop mode and another VSAN in the basic mode in the same switch without disrupting each other.
- Bring E ports out of isolation.

Zone Configuration

A zone can be configured using one of the following identifiers to assign members:

- pWWN—The WWN of the N or NL port in hex format (for example, 10:00:00:23:45:67:89:ab).
- Fabric port WWN—The WWN of the fabric port name in hex format (for example, 10:00:00:23:45:67:89:ab).
- FC ID—The N port ID in 0xhhhhhh format (for example, 0xce00d1).
- FC alias—The alias name is in alphabetic characters (for example, Payroll) and denotes a port ID or WWN. The alias can also include multiple members.
- Domain ID—The domain ID is an integer from 1 to 239. A mandatory port number of a non-Cisco switch is required to complete this membership configuration.

Send documentation comments to mdsfeedback-doc@cisco.com.

- **IP address**—The IP address of an attached device in 32 bytes in dotted decimal format along with an optional subnet mask. If a mask is specified, any device within the subnet becomes a member of the specified zone.
- **Interface**—Interface-based zoning is similar to port-based zoning because the switch interface is used to configure the zone. You can specify a switch interface as a zone member for both local and remote switches. To specify a remote switch, enter the remote switch WWN (sWWN) or the domain ID in the particular VSAN.

Configuring a Zone

To configure a zone and assign a zone name, follow these steps:

| | Command | Purpose |
|---------------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# zone name Zone1 vsan 3 switch(config-zone)# | Configures a zone called Zone 1 for the VSAN called vsan3. |
| Step 3 | switch(config-zone)# member <type> <value> pWWN example: switch(config-zone)# member pwwn 10:00:00:23:45:67:89:ab Fabric pWWN example: switch(config-zone)# member fwwn 10:01:10:01:10:ab:cd:ef FC ID example: switch(config-zone)# member fcid 0xce00d1 FC alias example: switch(config-zone)# member fcalias Payroll Domain ID example: switch(config-zone)# member domain-id 2 portnumber 23 FC alias example: switch(config-zone)# member ipaddress 10.15.0.0 255.255.0.0 Local sWWN interface example: switch(config-zone)# member interface fc 2/1 Remote sWWN interface example: switch(config-zone)# member interface fc2/1 swwn 20:00:00:05:30:00:4a:de Domain ID interface example: switch(config-zone)# member interface fc2/1 domain-id 25 | Configures a member for the specified zone (Zone1) based on the type (pWWN, fabric pWWN, FC ID, FC alias, domain ID, IP address, or interface) and value specified. |
| Tip | Use a relevant display command (for example, show interface or show flogi database) to obtain the required value in hex format. | |



Tip

Use the **show wwn switch** command to retrieve the sWWN. If you do not provide a sWWN, the software automatically uses the local sWWN.



Note

Interface-based zoning only works with Cisco MDS 9000 Family switches. Interface-based zoning does not work if interop mode is configured in that VSAN.

Send documentation comments to mdsfeedback-doc@cisco.com.

Alias Configuration

You can assign an alias name and configure an alias member using either the FC ID, fabric port WWN (fWWN), or pWWN values.



Tip As of Cisco MDS SAN-OS Release 1.3(4), the Cisco SAN-OS software supports a maximum of 2048 aliases per VSAN.

To create an alias, follow these steps:


| | Command | Purpose |
|--|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# fcalias name AliasSample vsan 3 switch(config-fcalias)# | Configures an alias name (AliasSample). |
| Step 3 | switch(config-fcalias)# member fcid 0x222222 | Configures alias members based on the specified FC ID type and value (0x222222). |
| | switch(config-fcalias)# member pwwn 10:00:00:23:45:67:89:ab | Configures alias members based on the specified port WWN type and value (pWWN 10:00:00:23:45:67:89:ab). |
| | switch(config-fcalias)# member fwwn 10:01:10:01:10:ab:cd:ef | Configures alias members based on the specified fWWN type and value (fWWN 10:01:10:01:10:ab:cd:ef). |
| Note Multiple members can be specified on multiple lines. | | |

[Send documentation comments to mdsfeedback-doc@cisco.com.](mailto:mdsfeedback-doc@cisco.com)

Zone Set Creation

In Figure 15-3, two separate sets are created, each with its own membership hierarchy and zone members.

Figure 15-3 Hierarchy of Zone Sets, Zones, and Zone Members



Zones provide a mechanism for specifying access control, while zone sets are a grouping of zones to enforce access control in the fabric. Either zone set A or zone set B can be activated (but not together).



Tip

Zone sets are configured with the names of the member zones and the VSAN (if the zone set is in a configured VSAN).

To create a zone set to include several zones, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# zoneset name Zoneset1 vsan 3 switch(config-zoneset) # | Configures a zone set called Zoneset1. Tip To activate a zone set, you must first create the zone and a zone set. |
| Step 3 | switch(config-zoneset) # member Zone1 | Adds Zone1 as a member of the specified zone set (Zoneset1). Tip If the specified zone name was not previously configured, this command will return the <code>Zone not present</code> error message. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|---|---|
| Step 4 | <pre>switch(config-zoneset)# zone name InlineZone1 switch(config-zoneset-zone)#</pre> | <p>Adds a zone (InlineZone1) to the specified zone set (Zoneset1).</p> <p>Tip Execute this step only if you need to create a zone from a zone set prompt.</p> |
| Step 5 | <pre>switch(config-zoneset-zone)# member fcid 0x111112 switch(config-zoneset-zone)#</pre> | <p>Adds a new member (FC ID 0x111112) to the newly created zone (InlineZone1).</p> <p>Tip Execute this step only if you need to add a member to a zone from a zone set prompt.</p> |

Active and Full Zone Set Considerations

Before configuring a zone set, consider the following guidelines:

- Each VSAN can have multiple zone sets but only one zone set can be active at any given time.
- When you create a zone set, that zone set becomes a part of the full zone set.
- When you activate a zone set, a copy of the zone set from the full zone set is used to enforce zoning, and is called the active zone set. An active zone set cannot be modified. A zone that is part of an active zone set is called an active zone.
- The administrator can modify the full zone set even if a zone set with the same name is active. However, the modification will be enforced only upon reactivation.
- When the activation is done, the active zone set is automatically stored in persistent configuration. This enables the switch to preserve the active zone set information across switch resets.
- All other switches in the fabric receive the active zone set so they can enforce zoning in their respective switches.
- Hard and soft zoning are implemented using the active zone set. Modifications take effect during zone set activation.
- An FC ID or Nx port that is not part of the active zone set belongs to the default zone and the default zone information is not distributed to other switches.



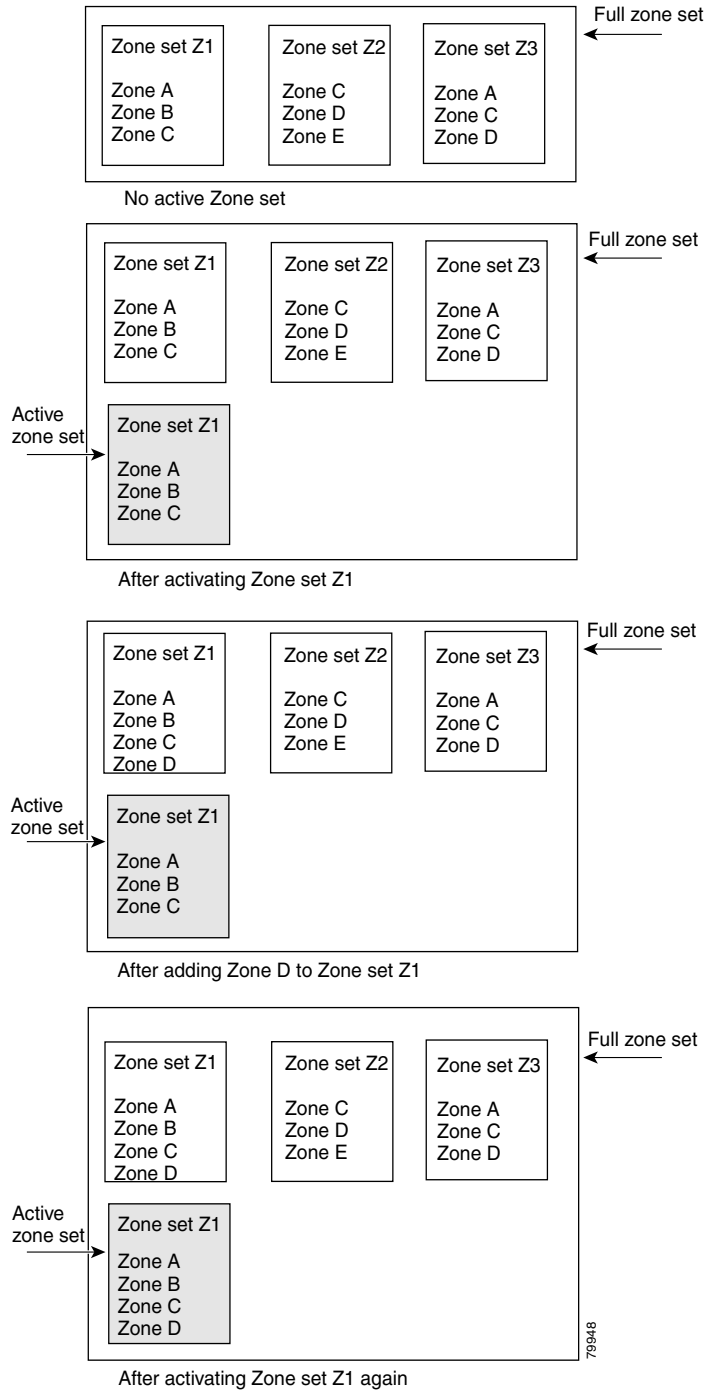
Note

If one zone set is active and you activate another zone set, the currently active zone set is automatically deactivated. You do not need to explicitly deactivate the currently active zone set before activating a new zone set.

Figure 15-4 shows a zone being added to an activated zone set.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 15-4 Active and Full Zone Sets



Send documentation comments to mdsfeedback-doc@cisco.com.

Activating a Zone Set

Changes to a zone set do not take effect to a full zone set until you activate it.
To activate a zone set, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# zoneset activate name Zoneset1 vsan 3 switch(config)# no zoneset activate name Zoneset1 vsan 3 | Activates the specified zone set. Deactivates the specified zone set |


Tip


You do not have to issue the **copy running-config startup-config** command to store the active zone set. However, you need to issue the **copy running-config startup-config** command to explicitly store full zone sets. It is not available across switch resets.

Zone Enforcement

Zoning can be enforced in two ways: soft and hard. Each end device (N port or NL port) discovers other devices in the fabric by querying the name server. When a device logs in to the name server, the name server returns the list of other devices that can be accessed by the querying device. If an Nx port does not know about the FC IDs of other devices outside its zone, it cannot access those devices.

In soft zoning, zoning restrictions are applied only during interaction between the name server and the end device. If an end device somehow knows the FC ID of a device outside its zone, it can access that device.

Hard zoning is enforced by the hardware on each frame sent by an Nx port. As frames enter the switch, source-destination IDs are compared with permitted combinations to allow the frame at wirespeed. Hard zoning is applied to all forms of zoning.


Note

Hard zoning enforces zoning restrictions on every frame, and prevents unauthorized access.

Switches in the Cisco MDS 9000 Family support both hard and soft zoning.

Send documentation comments to mdsfeedback-doc@cisco.com.

The Default Zone

Each member of a fabric (in effect a device attached to an Nx port) can belong to any zone. If a member is not part of any active zone, it is considered to be part of the default zone. Therefore, if no zone set is active in the fabric, all devices are considered to be in the default zone. Even though a member can belong to multiple zones, a member that is part of the default zone cannot be part of any other zone. The switch determines whether a port is a member of the default zone when the attached port comes up.



Note

Unlike configured zones, default zone information is not distributed to the other switches in the fabric.

Traffic can either be permitted or denied among members of the default zone. This information is not distributed to all switches; it must be configured in each switch.



Note

When the switch is initialized for the first time, no zones are configured and all members are considered to be part of the default zone. Members are not permitted to talk to each other.

Configure the default zone policy on each switch in the fabric. If you change the default zone policy on one switch in a fabric, be sure to change it on all the other switches in the fabric.



Note

The default settings for default zone configurations can be changed.

The default zone members are explicitly listed when the default policy is configured as permit or when a zone set is active. When the default policy is configured as deny, the members of this zone are not explicitly enumerated when you issue the **show zoneset active** command.

To permit or deny traffic in the default zone, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# zone default-zone permit vsan 1 | Permits traffic flow to default zone members. |
| | switch(config)# no zone default-zone permit vsan 1 | Denies traffic flow to default zone members and reverts to factory default. |

Zone Set Distribution

You can distribute full zone sets using one of two methods: at the EXEC mode level or at the configuration mode level. [Table 15-1](#) lists the differences.

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 15-1 zoneset distribution Command Differences

| zoneset distribute vsan Command (EXEC Mode) | zoneset distribute full vsan Command (Configuration Mode) |
|---|--|
| Distributes the full zone set immediately. | Does not distribute the full zone set immediately. |
| Does not distribute the full zone set information along with the active zone set during activation, deactivation, or merge process. | Remembers to distribute the full zone set information along with the active zone set during activation, deactivation, and merge processes. |

Enabling Full Zone Set Distribution

All switches in the Cisco MDS 9000 Family distribute active zone sets when new E port links come up or when a new zone set is activated in a VSAN. The zone set distribution takes effect while sending merge requests to the adjacent switch or while activating a zone set.

To enable full zone set and active zone set distribution to all switches on a per VSAN basis, follow these steps:

| | Command | Purpose |
|---------------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# zoneset distribute full vsan 33 | Enables sending a full zone set along with an active zone set. |

One-Time Distribution

As of Cisco MDS SAN-OS Release 1.3(4), you can perform a one-time distribution of inactive, unmodified zone sets throughout the fabric.

Use the **zoneset distribute vsan vsan-id** command in EXEC mode to perform this distribution.

```
switch# zoneset distribute vsan 2
Zoneset distribution initiated. check zone status
```

This command only distributes the full zone set information—it does not save the information to the startup configuration. You must explicitly issue the **copy running start** command to save the full zone set information to the startup configuration.



Note

The **zoneset distribute vsan vsan-id** command is supported in **interop 2** and **interop 3** modes—not in **interop 1** mode.

Send documentation comments to mdsfeedback-doc@cisco.com.

Use the **show zone status vsan *vsan-id*** command to check the status of the one-time zone set distribution request.

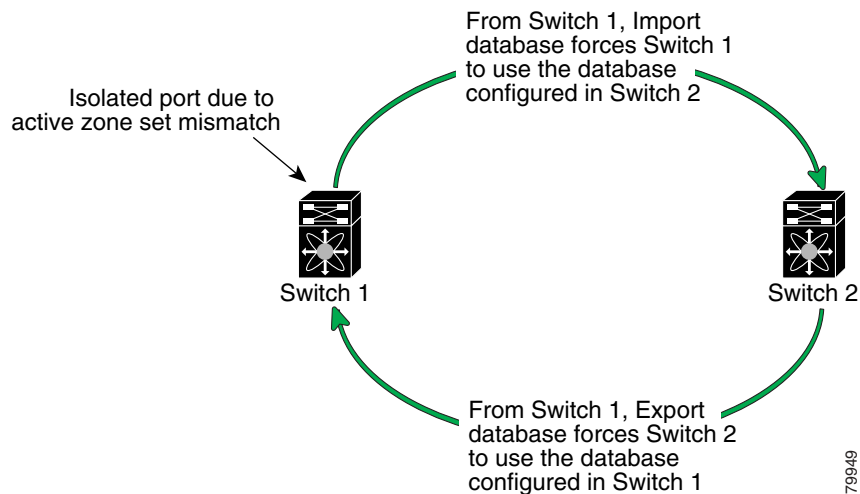
```
switch# show zone status vsan 2
VSAN: 3 default-zone: permit distribute: active only Interop: 100
      mode:basic merge-control:allow session:none
      hard-zoning:enabled
Default zone:
  qos:low broadcast:disabled ronly:disabled
Full Zoning Database :
  Zonesets:0 Zones:0 Aliases: 0
Active Zoning Database :
  Name: nozoneset Zonesets:1 Zones:2
Status: Zoneset distribution completed at 04:01:06 Aug 28 2004
```

Recovering from Link Isolation

When two switches in a fabric are merged using a TE or E port, these TE and E ports may become isolated when the active zone set databases are different between the two switches or fabrics. When a TE port or an E port become isolated, you can recover that port from its isolated state using one of three options:

- Import the neighboring switch's active zone set database and replace the current active zone set (see [Figure 15-5](#)).
- Export the current database to the neighboring switch (see [Figure 15-5](#)).
- Manually resolve the conflict by editing the full zone set, activating the corrected zone set, and then bringing up the link.

Figure 15-5 Importing and Exporting the Database



Send documentation comments to mdsfeedback-doc@cisco.com.

Importing and Exporting Zone Sets

To import or export the zone set information from or to an adjacent switch, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# zoneset import interface fc1/3 vsan 2 | Imports the zone set from the adjacent switch connected through the fc 1/3 interface for VSAN 2. |
| | switch# zoneset import interface fc1/3 vsan 2-5 | Imports the zone set from the adjacent switch connected through the fc 1/3 interface for VSANs ranging from 2 through 5. |
| Step 2 | switch# zoneset export vsan 5 | Exports the zone set to the adjacent switch connected through VSAN 5. |
| | switch# zoneset export vsan 5-8 | Exports the zone set to the adjacent switch connected through the range of VSANs 5 through 8. |



Note

Issue the **import** and **export** commands from a single switch. Importing from one switch and exporting from another switch can lead to isolation again.

Zone Set Duplication

You can make a copy and then edit it without altering the existing active zone set. You can copy an active zone set from the bootflash: directory, volatile: directory, or slot0, to one of the following areas:

- To the full zone set
- To a remote location (using FTP, SCP, SFTP, or TFTP).

The active zone set is not part of the full zone set. You cannot make changes to an existing zone set and activate it, if the full zone set is lost or is not propagated.



Caution

Copying an active zone set to a full zone set may overwrite a zone with the same name, if it already exists in the full zone set database.

Copying Zone Sets

On the Cisco MDS Family switches, you cannot edit an active zone set. However, you can copy an active zone set to create a new zone set that you can edit.

To copy zone sets, follow this step:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# zone copy active-zoneset full-zoneset vsan 2 Please enter yes to proceed.(y/n) [n]? y | Makes a copy of the active zone set in VSAN 2 to the full zone set. |
| | switch# zone copy vsan 3 active-zoneset scp://guest@myserver/tmp/active_zoneset.txt | Copies the active zone in VSAN 3 to a remote location using SCP. |

Send documentation comments to mdsfeedback-doc@cisco.com.

**Caution**

If the Inter-VSAN Routing (IVR) feature is enabled and if IVR zones exist in the active zone set, then a zone set copy operation copies all the IVR zones to the full zone database. To prevent copying to the IVR zones, you must explicitly remove them from the full zone set database before performing the copy operation. Refer to the [Chapter 17, “Configuring Inter-VSAN Routing”](#) for more information on the IVR feature.

Zone Database Information

If required, you can clear configured information stored in the zone server database.

**Note**

Clearing a zone set only erases the full zone database, not the active zone database.

Clearing the Zone Server Database

You can clear all configured information in the zone server database for the specified VSAN. To clear the zone server database, use the following command:

```
switch# clear zone database vsan 2
```

**Note**

After issuing a **clear zone database** command, you must explicitly issue the **copy running-config startup-config** to ensure that the running configuration is used when the switch reboots.

Zone-Based Traffic Priority

As of Cisco SAN-OS Release 2.0(1b), the zoning feature provides an additional segregation mechanism to prioritize select zones in a fabric and set up access control between devices. Using this feature, you can configure the Quality of Service (QoS) priority as a zone attribute. You can assign the QoS traffic priority attribute to be high, medium, or low. By default, zones with no specified priority are implicitly assigned a low priority. Zone-based QoS can only be implemented in Cisco MDS 9000 Family switches running Cisco MDS SAN-OS Release 2.0(1b) or later. See the [“VSAN Versus Zone-Based QoS” section on page 32-6](#) for more information.

To use this feature, you need to obtain the ENTERPRISE_PKG license (see [Chapter 3, “Obtaining and Installing Licenses”](#)) and you must enable QoS in the switch (see the [“QoS Initiation for Data Traffic” section on page 32-6](#)).

This feature allows SAN administrators to configure QoS in terms of a familiar data flow identification paradigm. You can configure this attribute on a zone-wide basis rather than between zone members.

**Caution**

If zone-based QoS is implemented in a switch, you cannot configure the interop mode in that VSAN.

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure zone priority, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# zone name QosZone vsan 2 switch(config-zone)# | Configures an alias name (QosZone). |
| Step 3 | switch(config-zone)# attribute qos priority high | Configures this zone to assign high priority QoS traffic to each frame matching this zone. |
| | switch(config-zone)# attribute qos priority medium | Configures this zone to assign medium priority QoS traffic to each frame matching this zone. |
| | switch(config-zone)# attribute qos priority low | Configures this zone to assign low priority QoS traffic to each frame matching this zone. |
| | switch(config-zone)# no attribute qos priority high | Reverts to using the default low priority for this zone. |
| Step 4 | switch(config)# zoneset name QosZoneset vsan 2 switch(config-zoneset)# | Configures a zone set called QosZone set. Tip To activate a zone set, you must first create the zone and a zone set. |
| Step 5 | switch(config-zoneset)# member QosZone switch(config-zoneset)# exit switch(config)# | Adds QosZone as a member of the specified zone set (QosZoneset). Tip If the specified zone name was not previously configured, this command will return the <code>Zone not present</code> error message. |
| Step 6 | switch(config)# zoneset activate name QosZoneset vsan 2 | Activates the specified zone set. |

Configuring Default Zone QoS Priority Attributes

QoS priority attribute configuration changes take effect when you activate the zone set of the associated zone.



Note

If a member is part of two zones with two different QoS priority attributes, the higher QoS value is implemented. This situation does not arise in the VSAN-based QoS as the first matching entry is implemented.

To configure the QoS priority attributes for a default zone, follow these steps:

| | Command | Purpose |
|--------|---|----------------------------------|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# zone default-zone vsan 1 switch(config-default-zone)# | Enters the default-zone submode. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|--|---|
| Step 3 | switch123(config-zone)# attribute qos priority high | Sets the QoS priority attribute for frames matching these zones. |
| | switch123(config-zone)# no attribute qos dscp | Removes the QoS priority attribute for the default zone and reverts to default low priority |

Configuring Broadcast Zoning

As of Release 2.0(1b), you can configure broadcast frames in the basic zoning mode. By default, broadcast zoning is disabled and broadcast frames are sent to all Nx ports in the VSAN. When enabled, broadcast frames are only sent to Nx ports in the same zone, or zones, as the sender. Enable broadcast zoning when a host or storage device uses this feature.



Tip

If any NL port attached to an FL port shares a broadcast zone with the source of the broadcast frame, then the frames are broadcast to all devices in the loop.



Caution

If broadcast zoning is enabled on a switch, you cannot configure the interop mode in that VSAN.

To broadcast frames in the basic zoning mode, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# zone broadcast enable vsan 2 | Broadcasts frames for the specified VSAN. |
| | switch(config)# no zone broadcast enable vsan 3 | Disables (default) broadcasting for the specified VSAN. |
| Step 3 | switch(config)# zone name BcastZone vsan 2 switch(config-zone)# | Creates a broadcast zone in the specified VSAN. |
| Step 4 | switch(config-zone)# member pwnn 21:00:00:20:37:f0:2e:4d | Adds the specified member to this zone. |
| Step 5 | switch(config-zone)# attribute broadcast | Specifies this zone to be broadcast to other devices. |
| Step 6 | switch(config-zone)# end switch# show zone vsan 2 zone name bcast-zone vsan 2 attribute broadcast pwnn 21:00:00:e0:8b:0b:66:56 pwnn 21:00:00:20:37:f0:2e:4d | Displays the broadcast configuration |

About LUN Zoning

Logical unit number (LUN) zoning is a feature specific to switches in the Cisco MDS 9000 Family.

Send documentation comments to mdsfeedback-doc@cisco.com.

**Caution**

LUN zoning can only be implemented in Cisco MDS 9000 Family switches. If LUN zoning is implemented in a switch, you cannot configure the interop mode in that switch.

**Note**

LUN zoning can be implemented in Cisco MDS 9000 Family switches running Cisco MDS SAN-OS Release 1.2 or later.

A storage device can have multiple LUNs behind it. If the device port is part of a zone, a member of the zone can access any LUN in the device. With LUN zoning, you can restrict access to specific LUNs associated with a device.

**Note**

When LUN 0 is not included within a zone, then, as per standards requirements, control traffic to LUN 0 (for example, REPORT_LUNS, INQUIRY) is supported, but data traffic to LUN 0 (for example, READ, WRITE) is denied.


- Host H1 can access LUN 2 in S1 and LUN 0 in S2. It cannot access any other LUNs in S1 or S2.
- Host H2 can access LUNs 1 and 3 in S1 and only LUN 1 in S2. It cannot access any other LUNs in S1 or S2.

**Note**

Unzoned LUNs automatically become members of the default zone.

Figure 15-6 shows a LUN-based zone example.

Figure 15-6 LUN Zoning Access



Send documentation comments to mdsfeedback-doc@cisco.com.

Configuring a LUN-Based Zone

To configure a LUN-based zone, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# zone name LunSample vsan 2 switch(config-zone)# | Configures a zone called LunSample for the specified VSAN (vsan 2). |
| Step 3 | switch(config-zone)# member pwnn 10:00:00:23:45:67:89:ab lun 64 | Configures a zone member based on the specified pWWN and LUN value. Note LUN x64 in hex format corresponds to 100 in decimal format. |
| | switch(config-zone)# member fcid 0x12465 lun 64 | Configures a zone member based on the FC ID and LUN value. |

Assigning LUNs to Storage Subsystems

LUN masking and mapping restricts server access to specific LUNs. If LUN masking is enabled on a storage subsystem and if you want to perform additional LUN zoning in a Cisco MDS 9000 Family switch, obtain the LUN number for each host bus adapter (HBA) from the storage subsystem and then configure the LUN-based zone procedure provided in the [“Configuring a LUN-Based Zone”](#) section on page 15-19.



Note

Refer to the relevant user manuals to obtain the LUN number for each HBA.



Caution

If you make any errors when configuring this scenario, you are prone to loose data.

About Read-Only Zones



Note

Read-only zoning can be implemented in Cisco MDS 9000 Family switches running Cisco MDS SAN-OS Release 1.2 or later.

By default, an initiator has both read and write access to the target's media when they are members of the same Fibre Channel zone. The read-only zone feature allows members to have only read access to the media within a read-only Fibre Channel zone.

You can also configure LUN zones as read-only zones.

Read-Only Zone Configuration Guidelines

Any zone can be identified as a read-only zone. By default all zones have read-write permission unless explicitly configured as a read-only zone.

Send documentation comments to mdsfeedback-doc@cisco.com.

Follow these guidelines when configuring read-only zones:

- If read-only zones are implemented, the switch prevents write access to user data within the zone.
- If two members belong to a read-only zone and to a read-write zone, read-only zone has priority and write access is denied.
- LUN zoning can only be implemented in Cisco MDS 9000 Family switches. If LUN zoning is implemented in a switch, you cannot configure interop mode in that switch.
- Read-only volumes are not supported by some operating system and file system combinations (for example, Windows NT or Windows 2000 and NTFS file system). Volumes within read-only zones are not available to such hosts. However, if these hosts are already booted when the read-only zones are activated, then read-only volumes are available to those hosts.

The read-only zone feature behaves as designed if FAT16 or FAT32 file system is used with the previously-mentioned Windows operating systems.

Configuring Read-Only Zones

To configure read-only zones, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# zone name Sample2 vsan 2 switch(config-zone)# | Configures a zone called Sample2 for the specified VSAN (vsan 2). |
| Step 3 | switch123(config-zone)# attribute read-only | Sets read-only attributes for the Sample2 zone. Note The default is read-write for all zones. |
| | switch123(config-zone)# no attribute read-only | Reverts the Sample2 zone attributes to read-write. |

To configure the **read-only** option for a default zone, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# zone default-zone vsan 1 switch(config-default-zone)# | Enters the default-zone submenu. |
| Step 3 | switch123(config-zone)# attribute read-only | Sets read-only attributes for the default zone. |
| | switch123(config-zone)# no attribute read-only | Reverts the default zone attributes to read-write (default). |

Send documentation comments to mdsfeedback-doc@cisco.com.

Renaming Zones, Zone Sets, fcaliases, and Zone Attribute Groups

To rename a zone, zone set, fcalias, or zone-attribute-group, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# zoneset rename oldname newname vsan 2 | Renames a zone set in the specified VSAN. |
| | switch(config)# zone rename oldname newname vsan 2 | Renames a zone in the specified VSAN. |
| | switch(config)# fcalias rename oldname newname vsan 2 | Renames a fcalias in the specified VSAN. |
| | switch(config)# zone-attribute-group rename oldname newname vsan 2 | Renames a zone attribute group in the specified VSAN. |
| Step 3 | switch(config)# zoneset activate name newname vsan 2 | Activates the zone set and updates the new zone name in the active zone set. |

Cloning Zones, Zone Sets, fcaliases, and Zone Attribute Groups

To clone a zone, zone set, fcalias, or zone-attribute-group, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# zoneset clone oldname newname vsan 2 | Clones a zone set in the specified VSAN. |
| | switch(config)# zone clone oldname newname vsan 2 | Clones a zone in the specified VSAN. |
| | switch(config)# fcalias clone oldname newname vsan 2 | Clones a fcalias in the specified VSAN. |
| | switch(config)# zone-attribute-group clone oldname newname vsan 2 | Clones a zone attribute group in the specified VSAN. |
| Step 3 | switch(config)# zoneset activate name newname vsan 2 | Activates the zone set and updates the new zone name in the active zone set. |

Displaying Zone Information

You can view any zone information by using the **show** command. If you request information for a specific object (for example, a specific zone, zone set, VSAN, or alias, or keywords such as **brief** or **active**), only information for the specified object is displayed. If you do not request specific information, all available information is displayed. See Examples 15-1 to 15-15.

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 15-1 Displays Zone Information for All VSANs

```
switch# show zone
zone name Zone3 vsan 1
  pwwn 21:00:00:20:37:6f:db:dd
  pwwn 21:00:00:20:37:9c:48:e5

zone name Zone2 vsan 2
  fwwn 20:41:00:05:30:00:2a:1e
  fwwn 20:42:00:05:30:00:2a:1e
  fwwn 20:43:00:05:30:00:2a:1e

zone name Zone1 vsan 1
  pwwn 21:00:00:20:37:6f:db:dd
  pwwn 21:00:00:20:37:a6:be:2f
  pwwn 21:00:00:20:37:9c:48:e5
  fcalias Alias1

zone name Techdocs vsan 3
  ip-address 10.15.0.0 255.255.255.0

zone name Zone21 vsan 5
  pwwn 21:00:00:20:37:a6:be:35
  pwwn 21:00:00:20:37:a6:be:39
  fcid 0xe000ef
  fcid 0xe000e0
  symbolic-nodename ign.test
  fwwn 20:1f:00:05:30:00:e5:c6
  fwwn 12:12:11:12:11:12:12:10
  interface fc1/5 swwn 20:00:00:05:30:00:2a:1e
  ip-address 12.2.4.5 255.255.255.0
  fcalias name Alias1 vsan 1
    pwwn 21:00:00:20:37:a6:be:35

zone name Zone2 vsan 11
  interface fc1/5 pwwn 20:4f:00:05:30:00:2a:1e

zone name Zone22 vsan 6
  fcalias name Alias1 vsan 1
    pwwn 21:00:00:20:37:a6:be:35

zone name Zone23 vsan 61
  pwwn 21:00:00:04:cf:fb:3e:7b lun 0000
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 15-2 Displays Zone Information for a Specific VSAN

```
switch# show zone vsan 1
zone name Zone3 vsan 1
  pwwn 21:00:00:20:37:6f:db:dd
  pwwn 21:00:00:20:37:9c:48:e5

zone name Zone2 vsan 1
  fwwn 20:4f:00:05:30:00:2a:1e
  fwwn 20:50:00:05:30:00:2a:1e
  fwwn 20:51:00:05:30:00:2a:1e
  fwwn 20:52:00:05:30:00:2a:1e
  fwwn 20:53:00:05:30:00:2a:1e

zone name Zone1 vsan 1
  pwwn 21:00:00:20:37:6f:db:dd
  pwwn 21:00:00:20:37:a6:be:2f
  pwwn 21:00:00:20:37:9c:48:e5
  fcalias Alias1
```

Use the **show zoneset** command to view the configured zone sets.

Example 15-3 Displays Configured Zone Set Information

```
switch# show zoneset vsan 1
zoneset name ZoneSet2 vsan 1
  zone name Zone2 vsan 1
    fwwn 20:4e:00:05:30:00:2a:1e
    fwwn 20:4f:00:05:30:00:2a:1e
    fwwn 20:50:00:05:30:00:2a:1e
    fwwn 20:51:00:05:30:00:2a:1e
    fwwn 20:52:00:05:30:00:2a:1e

  zone name Zone1 vsan 1
    pwwn 21:00:00:20:37:6f:db:dd
    pwwn 21:00:00:20:37:a6:be:2f
    pwwn 21:00:00:20:37:9c:48:e5
    fcalias Alias1

zoneset name ZoneSet1 vsan 1
  zone name Zone1 vsan 1
    pwwn 21:00:00:20:37:6f:db:dd
    pwwn 21:00:00:20:37:a6:be:2f
    pwwn 21:00:00:20:37:9c:48:e5
    fcalias Alias1
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 15-4 Displays Configured Zone Set Information for a Range of VSANs

```
switch# show zoneset vsan 2-3
zoneset name ZoneSet2 vsan 2
  zone name Zone2 vsan 2
    fwwn 20:52:00:05:30:00:2a:1e
    fwwn 20:53:00:05:30:00:2a:1e
    fwwn 20:54:00:05:30:00:2a:1e
    fwwn 20:55:00:05:30:00:2a:1e
    fwwn 20:56:00:05:30:00:2a:1e

  zone name Zone1 vsan 2
    pwwn 21:00:00:20:37:6f:db:dd
    pwwn 21:00:00:20:37:a6:be:2f
    pwwn 21:00:00:20:37:9c:48:e5
    fcalias Alias1

zoneset name ZoneSet3 vsan 3
  zone name Zone1 vsan 1
    pwwn 21:00:00:20:37:6f:db:dd
    pwwn 21:00:00:20:37:a6:be:2f
    pwwn 21:00:00:20:37:9c:48:e5
    fcalias Alias1
```

Use the **show zone name** command to display members of a specific zone.

Example 15-5 Displays Members of a Zone

```
switch# show zone name Zone1
zone name Zone1 vsan 1
  pwwn 21:00:00:20:37:6f:db:dd
  pwwn 21:00:00:20:37:a6:be:2f
  pwwn 21:00:00:20:37:9c:48:e5
  fcalias Alias1
```

Use the **show fcalias** command to display fcalias configuration.

Example 15-6 Displays fcalias Configuration

```
switch# show fcalias vsan 1
fcalias name Alias2 vsan 1

fcalias name Alias1 vsan 1
  pwwn 21:00:00:20:37:6f:db:dd
  pwwn 21:00:00:20:37:9c:48:e5
```

Use the **show zone member** command to display all zones to which a member belongs using the FC ID.

Example 15-7 Displays Membership Status

```
switch# show zone member pwwn 21:00:00:20:37:9c:48:e5
      VSAN: 1
zone Zone3
zone Zone1
fcalias Alias1
```

Use the **show zone statistics** command to display the number of control frames exchanged with other switches.

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 15-8 Displays Zone Statistics

```
switch# show zone statistics
Statistics For VSAN: 1
*****
Number of Merge Requests Sent: 24
Number of Merge Requests Recvd: 25
Number of Merge Accepts Sent: 25
Number of Merge Accepts Recvd: 25
Number of Merge Rejects Sent: 0
Number of Merge Rejects Recvd: 0
Number of Change Requests Sent: 0
Number of Change Requests Recvd: 0
Number of Change Rejects Sent: 0
Number of Change Rejects Recvd: 0
Number of GS Requests Recvd: 0
Number of GS Requests Rejected: 0
Statistics For VSAN: 2
*****
Number of Merge Requests Sent: 4
Number of Merge Requests Recvd: 4
Number of Merge Accepts Sent: 4
Number of Merge Accepts Recvd: 4
Number of Merge Rejects Sent: 0
Number of Merge Rejects Recvd: 0
Number of Change Requests Sent: 0
Number of Change Requests Recvd: 0
Number of Change Rejects Sent: 0
Number of Change Rejects Recvd: 0
Number of GS Requests Recvd: 0
Number of GS Requests Rejected: 0
```

Example 15-9 Displays LUN Zone Statistics

```
switch# show zone statistics lun-zoning
LUN zoning statistics for VSAN: 1
*****
S-ID: 0x123456, D-ID: 0x22222, LUN: 00:00:00:00:00:00:00
-----
Number of Inquiry commands received:          10
Number of Inquiry data No LU sent:            5
Number of Report LUNs commands received:      10
Number of Request Sense commands received:     1
Number of Other commands received:            0
Number of Illegal Request Check Condition sent: 0

S-ID: 0x123456, D-ID: 0x22222, LUN: 00:00:00:00:00:00:01
-----
Number of Inquiry commands received:          1
Number of Inquiry data No LU sent:            1
Number of Request Sense commands received:     1
Number of Other commands received:            0
Number of Illegal Request Check Condition sent: 0
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 15-10 Displays LUN Zone Statistics

```
switch# show zone statistics read-only-zoning
Read-only zoning statistics for VSAN: 2
*****
S-ID: 0x333333, D-ID: 0x111111, LUN: 00:00:00:00:00:00:64
-----
Number of Data Protect Check Condition Sent: 12
```

Example 15-11 Displays Active Zone Sets

```
switch# show zoneset active
zoneset name ZoneSet1 vsan 1
  zone name zone1 vsan 1
    fcid 0x080808
    fcid 0x090909
    fcid 0x0a0a0a
  zone name zone2 vsan 1
    * fcid 0xef0000 [pwwn 21:00:00:20:37:6f:db:dd]
    * fcid 0xef0100 [pwwn 21:00:00:20:37:a6:be:2f]
```

Example 15-12 Displays Brief Descriptions of Zone Sets

```
switch# show zoneset brief
zoneset name ZoneSet1 vsan 1
  zone zone1
  zone zone2
```

Example 15-13 Displays Active Zones

```
switch# show zone active
zone name Zone2 vsan 1
  * fcid 0x6c01ef [pwwn 21:00:00:20:37:9c:48:e5]

zone name IVRZ_IvrZone1 vsan 1
  pwwn 10:00:00:00:77:99:7a:1b
  * fcid 0xce0000 [pwwn 10:00:00:00:c9:2d:5a:dd]

zone name IVRZ_IvrZone4 vsan 1
  * fcid 0xce0000 [pwwn 10:00:00:00:c9:2d:5a:dd]
  * fcid 0x6c01ef [pwwn 21:00:00:20:37:9c:48:e5]

zone name Zone1 vsan 1667
  fcid 0x123456

zone name $default_zone$ vsan 1667
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 15-14 Displays Active Zone Sets

```
switch# show zoneset active
zoneset name ZoneSet4 vsan 1
  zone name Zone2 vsan 1
    * fcid 0x6c01ef [pwwn 21:00:00:20:37:9c:48:e5]

  zone name IVRZ_IvrZone1 vsan 1
    pwwn 10:00:00:00:77:99:7a:1b
    * fcid 0xce0000 [pwwn 10:00:00:00:c9:2d:5a:dd]

zoneset name QosZoneset vsan 2
  zone name QosZone vsan 2
  attribute qos priority high
  * fcid 0xce0000 [pwwn 10:00:00:00:c9:2d:5a:dd]
  * fcid 0x6c01ef [pwwn 21:00:00:20:37:9c:48:e5]

Active zoneset vsan 1667
  zone name Zone1 vsan 1667
    fcid 0x123456

  zone name $default_zone$ vsan 1667
```

Example 15-15 Displays Zone Status

```
switch# show zone status
VSAN: 1 default-zone: deny distribute: full Interop: Off
      mode:basic merge-control:allow session:none
      hard-zoning:enabled
Default zone:
  qos:low broadcast:disabled ronly:disabled
Full Zoning Database :
  Zonesets:1 Zones:11 Aliases:0
Active Zoning Database :
  Name: zoneset-1 Zonesets:1 Zones:11 Aliases:0
Status: Activation completed at Thu Feb 13 10:22:34 2003

VSAN: 2 default-zone: deny distribute: full Interop: Off
      mode:basic merge-control:allow session:none
      hard-zoning:enabled
Default zone:
  qos:low broadcast:disabled ronly:disabled
Full Zoning Database :
  Zonesets:1 Zones:10 Aliases:0
Active Zoning Database :
  Name: zoneset-2 Zonesets:1 Zones:10 Aliases:0
Status: Activation completed at Thu Feb 13 10:23:12 2003

VSAN: 3 default-zone: deny distribute: full Interop: Off
      mode:basic merge-control:allow session:none
      hard-zoning:enabled
Default zone:
  qos:low broadcast:disabled ronly:disabled
Full Zoning Database :
  Zonesets:1 Zones:10 Aliases:0
Active Zoning Database :
  Name: zoneset-3 Zonesets:1 Zones:10 Aliases:0
Status: Activation completed at Thu Feb 13 10:23:50 2003
```

Use the **show zone** command to display the zone attributes for all configured zones.

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 15-16 Displays Zone Statistics

```
switch# show zone
zone name lunSample vsan 1          <-----Read-write attribute
zone name ReadOnlyZone vsan 2      <-----Read-only attribute
      attribute read-only
```

Use the **show running** and **show zone active** commands to display the configured interface-based zones (see [Example 15-17](#) and [Example 15-18](#)).

Example 15-17 Displays the Interface-Based Zones

```
switch# show running
zone name if-zone vsan 1
      member interface fc2/15 swwn 20:00:00:0c:88:00:4a:e2
      member fwwn 20:4f:00:0c:88:00:4a:e2
      member interface fc2/1 swwn 20:00:00:05:30:00:4a:9e
      member pwwn 22:00:00:20:37:39:6b:dd
```

Example 15-18 Displays the fWWNs and Interfaces in an Active Zone

```
switch# show zone active
zone name if-zone vsan 1
  * fcid 0x7e00b3 [interface fc2/15 swwn 20:00:00:0c:88:00:4a:e2]
  * fcid 0x7e00b1 [interface fc2/15 swwn 20:00:00:0c:88:00:4a:e2]
  * fcid 0x7e00ac [interface fc2/15 swwn 20:00:00:0c:88:00:4a:e2]
  * fcid 0x7e00b3 [fwwn 20:4f:00:0c:88:00:4a:e2]
  * fcid 0x7e00b1 [fwwn 20:4f:00:0c:88:00:4a:e2]
  * fcid 0x7e00ac [fwwn 20:4f:00:0c:88:00:4a:e2]
      interface fc2/1 swwn 20:00:00:05:30:00:4a:9e
```

A similar output is also available on the remote switch (see [Example 15-19](#)).

Example 15-19 Displays the Local Interface Active Zone Details for a Remote Switch

```
switch# show zone active
zone name if-zone vsan 1
  * fcid 0x7e00b3 [interface fc2/15 swwn 20:00:00:0c:88:00:4a:e2]
  * fcid 0x7e00b1 [interface fc2/15 swwn 20:00:00:0c:88:00:4a:e2]
  * fcid 0x7e00ac [interface fc2/15 swwn 20:00:00:0c:88:00:4a:e2]
  * fcid 0x7e00b3 [fwwn 20:4f:00:0c:88:00:4a:e2]
  * fcid 0x7e00b1 [fwwn 20:4f:00:0c:88:00:4a:e2]
  * fcid 0x7e00ac [fwwn 20:4f:00:0c:88:00:4a:e2]
      interface fc2/1 swwn 20:00:00:05:30:00:4a:9e
```

About Enhanced Zoning

As of Cisco SAN-OS Release 2.0(1b), the zoning feature complies with the FC-GS-4 and FC-SW-3 standards. Both standards support the basic zoning functionalities explained in the previous section and the enhanced zoning functionalities described in this section.

Send documentation comments to mdsfeedback-doc@cisco.com.

Advantages of Enhanced Zoning

Table 15-2 lists the advantages of the enhanced zoning feature in all switches in the Cisco MDS 9000 Family.

Table 15-2 Advantages of Enhanced Zoning

| Basic Zoning | Enhanced Zoning | Enhanced Zoning Advantages |
|--|--|---|
| Administrators can make simultaneous configuration changes. Upon activation, one administrator can overwrite another administrator's changes. | Performs all configurations within a single configuration session. When you begin a session, the switch locks the entire fabric to implement the change. | One configuration session for the entire fabric to ensure consistency within the fabric. |
| If a zone is part of multiple zone sets, you create an instance of this zone in each zone set | References to the zone are used by the zone sets as required once you define the zone. | Reduced payload size as the zone is referenced. The size is more pronounced with bigger databases. |
| The default zone policy is defined per switch. To ensure smooth fabric operation, all switches in the fabric must have the same default zone setting. | Enforces and exchanges the default zone setting throughout the fabric. | Fabric-wide policy enforcement reduces troubleshooting time. |
| To retrieve the results of the activation on a per switch basis, the managing switch provides a combined status about the activation. It does not identify the failure switch. | Retrieves the activation results and the nature of the problem from each remote switch. | Enhanced error reporting eases the troubleshooting process |
| To distribute the zoning database, you must reactivate the same zone set. The reactivation may affect hardware changes for hard zoning on the local switch and on remote switches. | Implements changes to the zoning database and distributes it without reactivation. | Distribution of zone sets without activation avoids hardware changes for hard zoning in the switches. |
| The MDS-specific zone member types (IP address, symbolic node name, and other types) may be used by other non-Cisco switches. During a merge, the MDS-specific types can be misunderstood by the non-Cisco switches. | Provides a vendor ID along with a vendor-specific type value to uniquely identify a member type. | Unique vendor type. |
| The fWWN-based zone membership is only supported in Cisco interop mode. | Supports fWWN-based membership in the default interop mode. | The fWWN-based member type is standardized. |

Changing from Basic Zoning to Enhanced Zoning

To change to the enhanced zoning mode from the basic mode, follow these steps:

- Step 1** Verify that all switches in the fabric are capable of working in the enhanced mode.
If one or more switches are not capable of working in enhanced mode, then your request to move to enhanced mode is rejected.
- Step 2** Set the operation mode to enhanced zoning mode. By doing so, you will automatically start a session, acquire a fabric wide lock, distribute the active and full zoning database using the enhanced zoning data structures, distribute zoning policies and then release the lock. All switches in the fabric then move to the enhanced zoning mode.

Send documentation comments to mdsfeedback-doc@cisco.com.



Tip

After moving from basic zoning to enhanced zoning we recommend that you save the running configuration.

Changing from Enhanced Zoning to Basic Zoning

The standards do not allow you to move back to basic zoning. However, Cisco MDS switches allow this move to enable you to downgrade and upgrade to other Cisco SAN-OS releases.

To change to the basic zoning mode from the enhanced mode, follow these steps:

Step 1 Verify that the active and full zone set do not contain any configuration that is specific to the enhanced zoning mode.

If such configurations exist, delete them before proceeding with this procedure. If you do not delete the existing configuration, the Cisco SAN-OS software automatically removes them.

Step 2 Set the operation mode to basic zoning mode. By doing so, you will automatically start a session, acquire a fabric wide lock, distribute the zoning information using the basic zoning data structure, apply the configuration changes and release the lock from all switches in the fabric. All switches in the fabric then move to basic zoning mode.



Note

If a switch running Cisco MDS SAN-OS Release 2.0(1b), or later, with enhanced zoning enabled is downgraded to Cisco MDS SAN-OS Release 1.3(4), or earlier, the switch comes up in basic zoning mode and thus cannot join the fabric because all the other switches in the fabric are still in enhanced zoning mode.

Enabling Enhanced Zoning

By default, the enhanced zoning feature is disabled in all switches in the Cisco MDS 9000 Family.

To enable enhanced zoning in a VSAN, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# zone mode enhanced vsan 3000 Set zoning mode command initiated. Check zone status | Enables enhanced zoning in the specified VSAN. |
| | switch(config)# no zone mode enhanced vsan 150 Set zoning mode command initiated. Check zone status | Disables enhanced zoning in the specified. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Modifying the Zone Database

Modifications to the zone database is done within a session. A session is created at the time of the first successful configuration command. On creation of a session, a copy of the zone database is created. Any changes done within the session are performed on this copy of the zoning database. These changes in the copy zoning database are not applied to the effective zoning database, until you commit. the changes. Once you apply the changes, the session is closed.

If the fabric is locked by another user and for some reason the lock is not cleared, you can force the operation and close the session. You must have permission (role) to clear the lock in this switch and perform the operation on the switch from where the session was originally created.

To commit changes to the zoning database in a VSAN, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# conf t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# zone commit vsan 2 No pending info found | Applies the changes to the enhanced zone configuration and closes the session. |
| | switch(config)# no zone commit vsan 3 force | Forcefully applies the changes to the enhanced zone and closes the session created by another user. |
| | switch(config)# no zone commit vsan 2 No pending info found | Discards the changes made to the enhanced zone configuration. |
| | switch(config)# no zone commit vsan 3 force | Forcefully releases the lock and discards all changes performed by another user. |

Creating Attribute Groups

In enhanced mode, you can directly configure attributes using attribute groups.

To configure attribute groups, follow these steps:

-
- Step 1** Create an attribute group.
- ```
switch# conf t
switch(config)# zone-attribute-group name SampleAttributeGroup vsan 2
switch(config-attribute-group)#
```
- Step 2** Add the attribute to an attribute-group object.
- ```
switch(config-attribute-group)# readonly
switch(config-attribute-group)# broadcast
switch(config-attribute-group)# qos priority medium
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Step 3 Attach the attribute-group to a zone.

```
switch(config)# zone name Zone1 vsan 2
switch(config-zone)# attribute-group SampleAttributeGroup
switch(config-zone)# exit
switch(config)#
```

Step 4 Activate the zone set.

```
switch(config)# zoneset activate name Zoneset1 vsan 2
```

The attribute-groups are expanded and only the configured attributes are present in the active zone set.

Merging the Database

The merge behavior depends on the fabric-wide merge control setting:

- Restrict—If the two database are not identical, the ISLs between the switches are isolated.
- Allow—The two databases are merged using the merge rules specified in [Table 15-3](#).

Table 15-3 Database Zone Merge Status

| Local Database | Adjacent Database | Merge Status | Results of the Merge |
|--|-------------------|--------------|---|
| The databases contain zone sets with the same name ¹ but different zones, aliases, and attributes groups. | | Successful. | The union of the local and adjacent databases. |
| The databases contains a zone, zone alias, or zone attribute group object with same name ¹ but different members. | | Failed. | ISLs are isolated. |
| Empty. | Contains data. | Successful. | The adjacent database information populates the local database. |
| Contains data. | Empty. | Successful. | The local database information populates the adjacent database. |

1. In the enhanced zoning mode, the active zone set does not have a name in interop mode 1. The zone set names are only present for full zone sets.

The Merge Process

1. The software compares the protocol versions. If the protocol versions differ, then the ISL is isolated.
2. If the protocol versions are the same, then the zone policies are compared. If the zone policies differ, then the ISL is isolated.
3. If the zone merge options are the same, then the comparison is implemented based on the merge control setting.
 - a. If the setting is restrict, the active zone set and the full zone set should be identical. Otherwise the link is isolated.
 - b. If the setting is allow, then the merge rules are used to perform the merge (see [Table 15-3](#)).

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure merge control policies, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# zone merge-control restrict vsan 4 | Configures a restricted merge control setting for this VSAN. |
| | switch(config)# no zone merge-control restrict vsan 2 | Defaults to using the allow merge control setting for this VSAN. |
| | switch(config)# zone commit vsan 4 | Commits the changes made to VSAN 4. |

Default Zone Policies

To permit or deny traffic in the default zone, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# zone default-zone permit vsan 5 | Permits traffic flow to default zone members. |
| | switch(config)# no zone default-zone permit vsan 3 | Denies traffic flow to default zone members and reverts to factory default. |
| Step 3 | switch(config)# zone commit vsan 5 | Commits the changes made to VSAN 5. |

Broadcasting a Zone

You can specify an enhanced zone to restrict broadcast frames generated by a member in this zone to members within that zone. Use this feature when the host or storage devices support broadcasting.

Table 15-4 identifies the rules for the delivery of broadcast frames.

Table 15-4 *Broadcasting Requirements*

| Active Zoning? | Broadcast Enabled? | Frames Broadcast? | Comments |
|----------------|--------------------|-------------------|--|
| Yes | Yes | Yes | Broadcast to all Nx ports that share a broadcast zone with the source of broadcast frames. |
| No | Yes | Yes | Broadcast to all Nx ports. |
| Yes | No | No | Broadcasting is disabled. |



Tip

If any NL port attached to an FL port shares a broadcast zone with the source of the broadcast frame, then the frames are broadcast to all devices in the loop.

Send documentation comments to mdsfeedback-doc@cisco.com.

To broadcast frames in the enhanced zoning mode, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# confi t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# zone-attribute-group name BroadcastAttr vsan 2 | Configures the zone attribute group for the required VSAN. |
| | switch(config)# no zone-attribute-group name BroadcastAttr vsan 1 | Removes the zone attribute group for the required VSAN. |
| Step 3 | switch(config-attribute-group)# broadcast switch(config-attribute-group)# exit switch(config)# | Creates a broadcast attribute for this group and exits this submode. |
| | switch(config-attribute-group)# no broadcast | Removes broadcast attribute for this group and exits this submode. |
| Step 4 | switch(config)# zone name BroadcastAttr vsan 2 switch(config-zone)# | Configures a zone named BroadcastAttr in VSAN 2. |
| Step 5 | switch(config-zone)# member pwnn 21:00:00:e0:8b:0b:66:56 switch(config-zone)# member pwnn 21:01:00:e0:8b:2e:80:93 switch(config-zone)# attribute-group name BroadcastAttr switch(config-zone)# exit switch(config)# | Adds the specified members to this zone and exits this submode. |
| Step 6 | switch(config)# zone commit vsan 1 Commit operation initiated switch(config)# end | Applies the changes to the enhanced zone configuration and exits this submode. |
| Step 7 | switch# show zone vsan 1 zone name BroadcastAttr vsan 1 zone-attribute-group name BroadcastAttr vsan 1 broadcast pwnn 21:00:00:e0:8b:0b:66:56 pwnn 21:01:00:e0:8b:2e:80:93 | Displays the broadcast configuration |

Displaying Enhanced Zone Information

You can view any zone information by using the **show** command. See Examples 15-20 to 15-32.

Example 15-20 Displays the Active Zone Set Information for a Specified VSAN

```
switch# show zoneset active vsan 2
zoneset name testzoneset vsan 2
  zone name testzone vsan 2
    attribute read-only
    attribute broadcast
    attribute qos priority high
    pwnn 21:01:00:e0:8b:2e:a3:8a
    pwnn 22:00:00:0c:50:02:cb:59

zone name $default_zone$ vsan 2
  attribute read-only
  attribute qos priority high
  attribute broadcast]
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 15-21 Displays the e Zone Set Information for a Specified VSAN

```
switch# show zoneset vsan 2
zoneset name testzoneset vsan 2
  zone name testzone vsan 2
    zone-attribute-group name testattgp vsan 2
      read-only
      broadcast
      qos priority high
      pwwn 21:01:00:e0:8b:2e:a3:8a
      pwwn 22:00:00:0c:50:02:cb:59

zoneset name testzoneset2 vsan 2
  zone name testzone2 vsan 2
    pwwn 21:01:00:e0:8b:2e:68:8a
    pwwn 22:00:00:0c:50:02:cb:80

zoneset name testzoneset3 vsan 2
  zone name testzone3 vsan 2
    pwwn 21:01:00:e0:8b:2e:68:8a
    pwwn 22:00:00:0c:50:02:cb:80
```

Example 15-22 Displays the Zone Attribute Group Information for a Specified VSAN

```
switch# show zone-attribute-group vsan 2
zone-attribute-group name $default_zone_attr_group$ vsan 2
  read-only
  qos priority high
  broadcast
zone-attribute-group name testattgp vsan 2
  read-only
  broadcast
  qos priority high
```

Example 15-23 Displays the e fcalias Information for the Specified VSAN

```
switch# show fcalias vsan 2
fcalias name testfcalias vsan 2
  pwwn 21:00:00:20:37:39:b0:f4
  pwwn 21:00:00:20:37:6f:db:dd
  pwwn 21:00:00:20:37:a6:be:2f
```

Example 15-24 Displays the Zone Status for the Specified VSAN

```
switch# show zone status vsan 2
VSAN: 2 default-zone: permit distribute: active only Interop: 100
  mode:basic merge-control:allow session:none
  hard-zoning:enabled
Default zone:
  qos:low broadcast:disabled ronly:disabled
Full Zoning Database :
  Zonesets:3 Zones:3 Aliases: 0 Attribute-groups: 2
Active Zoning Database :
  Name: testzoneset Zonesets:1 Zones:2
Status:
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 15-25 Displays an Active Zone Status for the Specified VSAN

```
switch# show zone status vsan 1
VSAN: 1 default-zone: permit distribute: full Interop: 100
      mode: enhanced merge-control: allow session: active <-----Indicates an active session.
      Hard zoning is enabled
Default zone:
      qos:low broadcast:disabled ronly:disabled
Full Zoning Database :
      Zonesets:4 Zones:4 Aliases: 0 Attribute-groups: 1
Active Zoning Database :
      Database Not Available
Status: Set zoning mode complete at 10:36:48 Aug 18 2004
```

Example 15-26 Displays the Pending Zone Set Information for the VSAN to be Committed

```
switch# show zoneset pending vsan 2
No pending info found
```

Example 15-27 Displays the Pending Zone Information for the VSAN to be Committed

```
switch# show zone pending vsan 2
No pending info found
```

Example 15-28 Displays the Pending Zone Information for the VSAN to be Committed

```
switch# show zone-attribute-group pending vsan 2
No pending info found
```

Example 15-29 Displays the Pending Active Zone Set Information for the VSAN to be Committed

```
switch# show zoneset pending active vsan 2
No pending info found
```

Example 15-30 Displays the Difference between the Pending and Effective Zone Information for the Specified VSAN

```
switch# show zone pending-diff vsan 2
zone name testzone vsan 2
- member pwnn 21:00:00:20:37:4b:00:a2
+ member pwnn 21:00:00:20:37:60:43:0c
```

Exchange Switch Support (ESS) defines a mechanism for two switches to exchange various supported features (see [Example 15-30](#)).

Example 15-31 Displays the ESS Information for All Switches in the Specified VSAN

```
switch# show zone ess vsan 2
ESS info on VSAN 2 :
  Domain : 210, SWWN : 20:02:00:05:30:00:85:1f, Cap1 : 0xf3, Cap2 : 0x0
```

Example 15-32 Displays the Pending fcalias Information for the VSAN to be Committed

```
switch# show fcalias pending vsan 2
No pending info found
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Default Settings

Table 15-5 lists the default settings for basic zone parameters.

Table 15-5 ***Default Basic Zone Parameters***

| Parameters | Default |
|-----------------------------|--|
| Default zone policy | Denied to all members. |
| Full zone set distribute | The full zone set(s) is not distributed. |
| Zone based traffic priority | Low. |
| Read-only zones | Read-write attributes for all zones. |
| Broadcast frames | Sent to all Nx ports. |
| Broadcast zoning | Disabled. |
| Enhanced zoning | Disabled. |

Send documentation comments to mdsfeedback-doc@cisco.com.



Distributing Device Alias Services

As of Release 2.0(1b), all switches in the Cisco MDS 9000 Family offer a new alias distribution feature called Distributed Device Alias Services (device alias). In Release 1.3 and earlier, aliases were distributed on a per VSAN basis. Using this new, enhanced service, you now have the option to distribute device alias names on a fabric-wide basis.

This chapter includes the following sections:

- [About Device Aliases, page 16-1](#)
- [Device Alias Features, page 16-2](#)
- [Device Alias Requirements, page 16-2](#)
- [Zone Aliases Versus Device Aliases, page 16-2](#)
- [Modifying the Device Alias Database, page 16-3](#)
- [Fabric Lock Override, page 16-4](#)
- [Device Alias Distribution, page 16-5](#)
- [Legacy Zone Alias Configuration Conversion, page 16-5](#)
- [“Database Merge Guidelines” section on page 16-5](#)
- [Device Alias Statistics Cleanup, page 16-6](#)
- [Device Alias Configuration Verification, page 16-6](#)
- [Default Settings, page 16-10](#)

About Device Aliases

When the port WWN of a device must be specified to configure different features (zoning, QoS, port security) in a Cisco MDS 9000 Family switch, you must assign the right device name each time you configure these features. An inaccurate device name may cause unexpected results. You can circumvent this problem if you define a user-friendly name for a port WWN and use this name in all the configuration commands as required. These user friendly names are referred to as *device aliases* in this chapter.

As of Release 2.0(1b), all switches in the Cisco MDS 9000 Family offer a new alias distribution feature called Distributed Device Alias Services. In Release 1.3 and earlier, device aliases were distributed on a per VSAN basis. Using this new, enhanced service, you now have the option to distribute device alias names on a fabric-wide basis. Device alias distribution allows you to move host bus adapters (HBAs) between VSANs without manually reentering alias names.

[Send documentation comments to mdsfeedback-doc@cisco.com.](mailto:mdsfeedback-doc@cisco.com)

Device Alias Features

Device aliases have the following features:

- The device alias information is independent of your VSAN configuration.
- The device alias configuration and distribution is independent of the zone server and the zone server database.
- You can import legacy zone alias configurations without losing data.
- The device alias application uses the Cisco Fabric Services (CFS) infrastructure to enable efficient database management and distribution. Device aliases use the coordinated distribution mode and the fabric-wide distribution scope (see [Chapter 9, “Using the CFS Infrastructure”](#)).
- When you configure zones, IVR zones, or QoS features using device aliases, and if you display these configuration, you will automatically see that the device aliases are displayed along with their respective pWWNs.

Device Alias Requirements

Device aliases have the following requirements:

- You can only assign device aliases to pWWNs.
- Ensure that the mapping between the pwwn and the device alias to which it is mapped has a one to one relationship. A pWWN can be mapped to only one device alias and vice versa.
- A device alias name is restricted to 64 alphanumeric characters and may include one or more of the following characters:
 - a to z and A to Z
 - 1 to 9
 - - (hyphen) and _ (underscore)
 - \$ and ^

Zone Aliases Versus Device Aliases

[Table 16-1](#) compares the configuration differences between zone-based alias configuration and device alias configuration.

Table 16-1 *Comparison Between Zone Aliases and Device Aliases*

| Zone-Based Aliases | Device Aliases |
|--|---|
| Aliases are limited to the specified VSAN | You can define device aliases without specifying the VSAN number. You can also use the same definition in one or more VSANs without any restrictions. |
| Zone aliases are part of the zoning configuration, the alias mapping cannot be used to configure other features. | Device aliases can be used to any feature that uses the pWWN. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 16-1 **Comparison Between Zone Aliases and Device Aliases (continued)**

| Zone-Based Aliases | Device Aliases |
|--|---|
| You can use any zone member type to specify the end devices. | Only pWWNs are supported along with new device aliases like IP addresses. |
| Configuration is contained within the Zone Server database and is not available to other features. | Device aliases are not restricted to zoning. Device alias configuration is available to the FCNS, zone, fcping, traceroute, and IVR applications. |

Modifying the Device Alias Database

The device alias feature uses two databases to accept and implement device alias configurations.

- Effective database—The database currently used by the fabric.
- Pending database—Your subsequent device alias configuration changes are stored in the pending database.

If you modify the device alias configuration, you need to commit or discard the changes as the fabric remains locked during this period.

Locking The Fabric

When you perform the first device alias task (regardless of which device alias task), the fabric is automatically locked for the device alias feature. Once you lock the fabric, the following situations apply:

- No other user can make any configuration changes to this feature.
- A copy of the effective database is obtained and used as the pending database. Modifications from this point on are made to the pending database. The pending database remains in effect until you commit the modifications to the pending database or discard (**abort**) the changes to the pending database.

To lock the fabric and modify the device alias configuration in the pending database, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# device-alias database switch(config-device-alias-db)# | Enters the pending database configuration submode. |
| Step 3 | switch(config-device-alias-db)# device-alias name x pwwn 21:01:00:e0:8b:2e:80:93 | Specifies a device name (x) for the device that is identified by its pWWN. Starts writing to the pending database and simultaneously locks the fabric as this is the first-issued device alias configuration command. |
| | switch(config-device-alias-db)# no device-alias name Doc | Removes the device name (SampleName) for the device that is identified by its pWWN. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Committing Changes

If you commit the changes made to the pending database, the following events occur:

1. The pending database contents overwrites the effective database contents.
2. The pending database is emptied of its contents.
3. The fabric lock is released for this feature.

To commit the changes, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# device-alias commit | Commits the changes made to the currently active session. |

Discarding Changes

If you discard the changes made to the pending database, the following events occur:

1. The effective database contents remain unaffected.
2. The pending database is emptied of its contents.
3. The fabric lock is released for this feature.

To discard the device alias session, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# device-alias abort | Discards the currently active session. |

Fabric Lock Override

If you have performed a device alias task and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.



Tip

The changes are only available in the volatile directory and are subject to being discarded if the switch is restarted.

To use administrative privileges and release a locked device alias session, use the **clear device-name session** command in EXEC mode.

```
switch# clear device-alias session
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Device Alias Distribution

By default, device alias distribution is enabled. The device alias feature uses the coordinated distribution mechanism to distribute the modifications to all switches in a fabric.

If you have not committed the changes and you enable distribution, then a commit task will fail (see [Example 16-10](#)).

To enable the device alias distribution, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# device-alias distribute switch(config)# no device-alias distribute | Enables the distribution (default). Disables the distribution. |

Legacy Zone Alias Configuration Conversion

You can import legacy zone alias configurations to use this feature without losing data, if they satisfy the following restrictions:

- Each zone alias has only one member.
- The member type is pWWN.
- The name and definition of the zone alias should not be the same as any existing device alias name.

If any name conflict exists, the zone aliases are not imported.



Tip

Ensure to copy any required zone aliases to the device alias database as required by your configuration.

When an import operation is complete, the modified alias database is distributed to all other switches in the physical fabric when you perform the **commit** operation. At this time if you do not want to distribute the configuration to other switches in the fabric, you can perform the **abort** operation and the merge changes are completely discarded.

To import the zone alias for a specific VSAN, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# device-alias import fcalias vsan 3 | Imports the fcalias information for the specified VSAN. |

Database Merge Guidelines

Refer to the “[CFS Merge Support](#)” section on [page 9-7](#) for detailed concepts.

Send documentation comments to mdsfeedback-doc@cisco.com.

When merging two device alias databases, follow these guidelines:

- Verify that two device aliases with different names are not mapped to the same pWWN.
- Verify that two identical pWWNs are not mapped to two different device aliases.
- Verify that the combined number of the device aliases in both databases does not exceed 8191 (8K). For example, if Database N has 6000 device aliases and Database M has 2192 device aliases, this merge operation will fail.

Device Alias Statistics Cleanup

Use the **clear device-name statistics** command to clear device alias statistics (for debugging purposes):

```
switch# clear device-alias statistics
```

Device Alias Configuration Verification

You can view device alias information by using the **show device-alias** command. See Examples 16-1 to 16-20.

Example 16-1 Displays All Configured Device Aliases from the Effective Database

```
switch# show device-alias database
device-alias name SampleName pwwn 21:00:00:e0:8b:0b:66:56
device-alias name x pwwn 21:01:00:e0:8b:2e:80:93
```

Total number of entries = 2

Example 16-2 Displays the Specified Device Name

```
switch# show device-alias name x
device-alias name x pwwn 21:01:00:e0:8b:2e:80:93
```

Example 16-3 Displays the Pending Database with No Modifications

```
switch# show device-alias database pending
There are no pending changes
```

Example 16-4 Displays the Pending Database with Modifications

```
switch# show device-alias database pending
device-alias name x pwwn 21:01:00:e0:8b:2e:80:93
device-alias name SampleName pwwn 21:00:00:e0:8b:0b:66:56
device-alias name y pwwn 21:00:00:20:37:39:ab:5f
device-alias name z pwwn 21:00:00:20:37:39:ac:0d
```

Total number of entries = 4

Example 16-5 Displays the Specified Device Name in the Pending Database

```
switch# show device-alias name x pending
device-alias name x pwwn 21:01:00:e0:8b:2e:80:93
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 16-6 Displays the Specified pWWN in the Pending Database

```
switch# show device-alias pwwn 21:01:00:e0:8b:2e:80:93 pending
device-alias name x pwwn 21:01:00:e0:8b:2e:80:93
```

Example 16-7 Displays the Difference between the Pending and Effective Databases

```
switch# show device-alias database pending-diff
- device-alias name Doc pwwn 21:01:02:03:00:01:01:01
+ device-alias name SampleName pwwn 21:00:00:e0:8b:0b:66:56
```

Example 16-8 Displays the Specified pWWN

```
switch# show device-alias pwwn 21:01:01:01:01:11:01:01
device-alias name Doc pwwn 21:01:01:01:01:11:01:01
```

Example 16-9 Displays a Successful Device Alias Status

```
switch# show device-alias status
Fabric Distribution: Enabled <-----Distribution is enabled
Database:-Device Aliases 24
Locked By:-User "Test" SWWN 20:00:00:0c:cf:f4:02:83<-Lock holder's user name and switch ID
Pending Database:- Device Aliases 24
Status of the last CFS operation issued from this switch:
=====
Operation: Enable Fabric Distribution
Status: Success
```

Example 16-10 Displays a Failed Device Alias Status

```
switch# show device-alias status
Fabric Distribution: Disabled
Database:- Device Aliases 25
Status of the last CFS operation issued from this switch:
=====
Operation: Commit
Status: Failed (Reason: Operation is not permitted as the fabric distribution is
currently disabled.)
```

Example 16-11 Displays the Device Alias Status of a abort Command

```
switch# show device-alias status
Fabric Distribution: Enabled
Database:- Device Aliases 24
Status of the last CFS operation issued from this switch:
=====
Operation: Abort
Status: Success
```

Example 16-12 Displays the Device Alias Status of a Cleared Session

```
switch# show device-alias status
Fabric Distribution: Enabled
Database:- Device Aliases 24
Status of the last CFS operation issued from this switch:
=====
Operation: Clear Session <-----Lock released by administrator
Status: Success <-----Successful status of the operation
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 16-13 Displays the Device Alias Status When Distribution Is Disabled

```
switch# show device-alias status
Fabric Distribution: Disabled
Database:- Device Aliases 24
Status of the last CFS operation issued from this switch:
=====
Operation: Disable Fabric Distribution
Status: Success
```

Example 16-14 Displays the Device Alias in the FLOGI Database

```
switch# show flogi database
-----
INTERFACE  VSAN    FCID          PORT NAME          NODE NAME
-----
fc2/9      1       0x670100     21:01:00:e0:8b:2e:80:93  20:01:00:e0:8b:2e:80:93
[x] <-----Device alias name
fc2/12     1       0x670200     21:00:00:e0:8b:0b:66:56  20:00:00:e0:8b:0b:66:56
[SampleName] <-----Device alias name

Total number of flogi = 2
```

Example 16-15 Displays the Device Alias in the FCNS Database

```
switch# show fcns database

VSAN 1:
-----
FCID      TYPE  PWWN          (VENDOR)          FC4-TYPE:FEATURE
-----
0x670100  N     21:01:00:e0:8b:2e:80:93 (Qlogic)          scsi-fcp:init
[x]
0x670200  N     21:00:00:e0:8b:0b:66:56 (Qlogic)          scsi-fcp:init
[SampleName]

Total number of entries = 2
```

Example 16-16 Displays the fcping Statistics for the Specified Device Alias

```
switch# fcping device-alias x vsan 1
28 bytes from 21:01:00:e0:8b:2e:80:93 time = 358 usec
28 bytes from 21:01:00:e0:8b:2e:80:93 time = 226 usec
28 bytes from 21:01:00:e0:8b:2e:80:93 time = 372 usec
```

Example 16-17 Displays the fctrace Information for the Specified Device Alias

```
switch# fctrace device-alias x vsan 1
Route present for : 21:01:00:e0:8b:2e:80:93
20:00:00:05:30:00:4a:e2(0xfffc67)
```

Where available, device aliases are displayed regardless of a member being configured using a **device-alias** command or a zone-specific **member pwwn** command (see [Example 16-18](#) and [Example 16-19](#)).

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 16-18 Displays the Device Aliases in the Zone Set Information

```
switch# show zoneset
zoneset name s1 vsan 1
  zone name z1 vsan 1
    pwwn 21:01:00:e0:8b:2e:80:93 [x] <-----Device alias displayed for each pWWN.
    pwwn 21:00:00:20:37:39:ab:5f [y]

  zone name z2 vsan 1
    pwwn 21:00:00:e0:8b:0b:66:56 [SampleName]
    pwwn 21:00:00:20:37:39:ac:0d [z]
```

Example 16-19 Displays the the Device Aliases in the Active Zone Set

```
switch# show zoneset active
zoneset name s1 vsan 1
  zone name z1 vsan 1
    * fcid 0x670100 [pwwn 21:01:00:e0:8b:2e:80:93] [x]
    pwwn 21:00:00:20:37:39:ab:5f [y]

  zone name z2 vsan 1
    * fcid 0x670200 [pwwn 21:00:00:e0:8b:0b:66:56] [SampleName]
    pwwn 21:00:00:20:37:39:ac:0d [z]
```

Example 16-20 Displays Statistics for the Device Alias Application

```
switch# show device-alias statistics
      Device Alias Statistics
=====
Lock requests sent: 2
Database update requests sent: 1
Unlock requests sent: 1
Lock requests received: 1
Database update requests received: 1
Unlock requests received: 1
Lock rejects sent: 0
Database update rejects sent: 0
Unlock rejects sent: 0
Lock rejects received: 0
Database update rejects received: 0
Unlock rejects received: 0
Merge requests received: 0
Merge request rejects sent: 0
Merge responses received: 2
Merge response rejects sent: 0
Activation requests received: 0
Activation request rejects sent: 0
Activation requests sent: 2
Activation request rejects received: 0
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Default Settings

Table 16-2 lists the default settings for device alias parameters.

Table 16-2 Default Device Alias Parameters

| Parameters | Default |
|--------------------------------|--|
| Database in use | Effective database. |
| Database to accept changes | Pending database. |
| Device alias fabric lock state | Locked with the first device alias task. |



Configuring Inter-VSAN Routing

This chapter explains the Inter-VSAN Routing (IVR) feature and provides details on sharing resources across VSANs using IVR management interfaces provided in the switch.

This chapter includes the following sections:

- [About IVR, page 17-2](#)
- [IVR Terminology, page 17-3](#)
- [IVR Guidelines, page 17-4](#)
- [IVR Configuration, page 17-5](#)
- [Unique Domain ID Configuration Options, page 17-6](#)
- [Enabling IVR, page 17-6](#)
- [Database Merge Guidelines, page 17-23](#)
- [About IVR Topologies, page 17-10](#)
- [Configuring IVR Topologies, page 17-10](#)
- [Adding IVR Virtual Domain, page 17-15](#)
- [About IVZs and IVZSs, page 17-16](#)
- [Configuring IVZs and IVZSs, page 17-18](#)
- [IVR Interoperability, page 17-22](#)
- [Configuring IVR Using Read-Only Zoning, page 17-22](#)
- [Configuring IVR Logging Levels, page 17-25](#)
- [Displaying IVR Information, page 17-25](#)
- [Sample Configuration, page 17-28](#)
- [Default Settings, page 17-31](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

About IVR

Virtual SANs (VSANs) improve storage area network (SAN) scalability, availability, and security by allowing multiple Fibre Channel SANs to share a common physical infrastructure of switches and ISLs. These benefits are derived from the separation of Fibre Channel services in each VSAN and isolation of traffic between VSANs. Data traffic isolation between the VSANs also inherently prevents sharing of resources attached to a VSAN, such as robotic tape libraries. Using IVR, you can access resources across VSANs without compromising other VSAN benefits.

Data traffic is transported between specific initiators and targets on different VSANs without merging VSANs into a single logical fabric. Fibre Channel control traffic does not flow between VSANs, and initiators cannot access any resource across VSANs other than the resources designated.

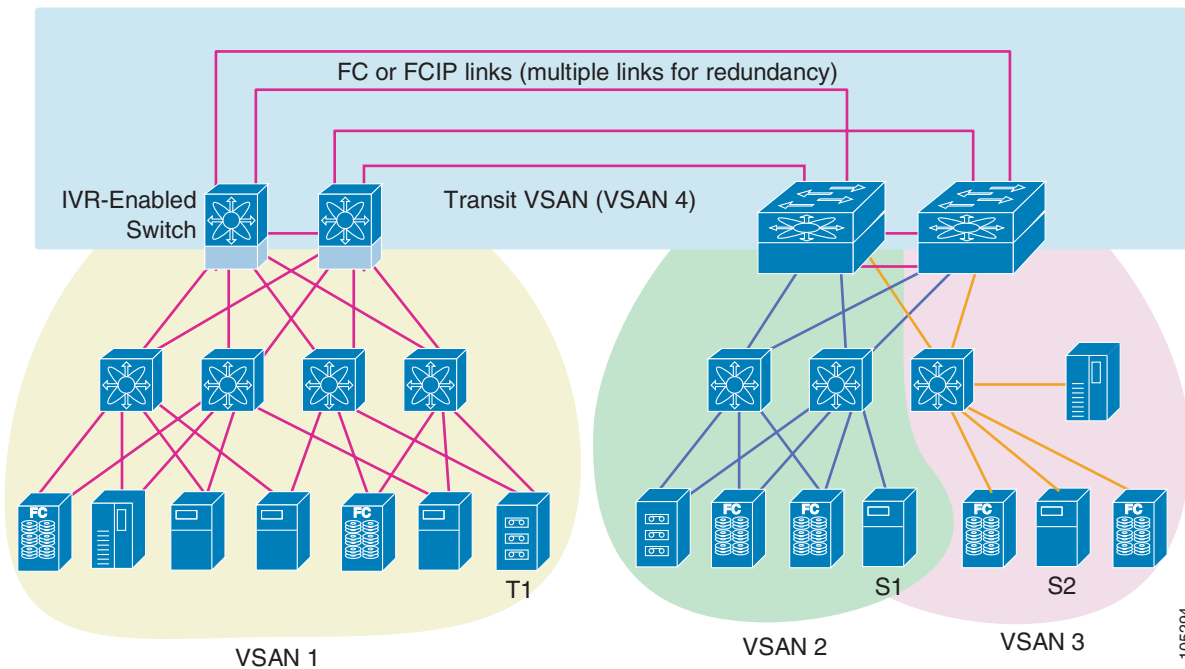
IVR is not limited to VSANs present on a common switch. Routes that traverse one or more VSANs across multiple switches can be established, if necessary, to establish proper interconnections. IVR used in conjunction with FCIP provides more efficient business continuity or disaster recovery solutions (see [Figure 17-1](#)).



Note

See the “[Sample Configuration](#)” section on [page 17-28](#) for procedures to configure the sample scenario shown in [Figure 17-1](#).

Figure 17-1 Traffic Continuity Using IVR and FCIP



105294

Send documentation comments to mdsfeedback-doc@cisco.com.

IVR Features

IVR supports the following features:

- Accesses resources across VSANs without compromising other VSAN benefits.
- Transports data traffic between specific initiators and targets on different VSANs without merging VSANs into a single logical fabric.
- Shares valuable resources (like tape libraries) across VSANs without compromise.
- Provides efficient business continuity or disaster recovery solutions when used in conjunction with FCIP.
- Is in compliance with Fibre Channel standards.
- Incorporates third-party switches, however, IVR-enabled VSANs may have to be configured in one of the interop modes.

IVR Terminology

The following IVR-related terms are used in this chapter:

- Native VSAN—The VSAN to which an end device logs on is the native VSAN for that end device.
- Inter-VSAN zone (IVZ)—A set of end devices that are allowed to communicate across VSANs within their interconnected SAN fabric. This definition is based on their port world wide names (pWWNs) and their native VSAN associations. As of Cisco MDS SAN-OS Release 1.3(1b), you can configure up to 200 IVZs and 2000 IVZ members on the switches in the network. As of Cisco MDS SAN-OS Release 2.1(1a), you can configure up to 2000 IVZs and 10,000 IVZ members on the switches in the network.
- Inter-VSAN zone sets (IVZS)—One or more IVZs make up an IVZS. You can configure up to 32 IVZSs on any switch in the Cisco MDS 9000 Family. Only one IVZS can be active at any time.
- IVR path—An IVR path is a set of switches and Inter-Switch Links through which a frame from an end device in one VSAN can reach another end device in another VSAN. Multiple paths can exist between two end devices.
- IVR-enabled switch—A switch on which the IVR feature is enabled.
- Edge VSAN—A VSAN that initiates (source edge-VSAN) or terminates (destination edge-VSAN) an IVR path. Edge VSANs can be adjacent to each other or they can be connected by one or more transit VSANs. In [Figure 17-1](#), VSANs 1, 2, and 3 are edge VSANs.

**Note**

An edge VSAN for one IVR path can be a transit VSAN for another IVR path.

- Transit VSAN—A VSAN that exists along an IVR path from a source edge VSAN to the destination edge VSAN. In [Figure 17-1](#), VSAN 4 is a transit VSAN.

**Note**

When the source and destination edge VSANs are adjacent to each other, then a transit VSAN is not required between them.

Send documentation comments to mdsfeedback-doc@cisco.com.

- Border switch—An IVR-enabled switch that is a member of two or more VSANs, such as the IVR-enabled switch between VSAN 1 and VSAN 4 in [Figure 17-1](#).
- Edge switch—A switch to which a member of an IVR zone has logged in. Edge switches are unaware of the IVR configurations in the border switches. Edge switches need not be IVR enabled.

IVR Guidelines

Before configuring an IVR SAN fabric, consider the following guidelines:

- Configure unique domain IDs across all VSANs and switches participating in IVR operations. The following switches participate in IVR operations:
 - All edge switches in the edge VSANs (source and destination)
 - All switches in transit VSANs
- Configure IVR only in the relevant border switches.
- Acquire a mandatory Enterprise License Package or, as of Cisco MDS SAN-OS Release 2.1(1a), a SAN Extension over IP License package for this feature.



Tip

If you change any FSPF link cost, ensure that the FSPF path distance (that is, the sum of the link costs on the path) of any IVR path is less than 30,000.



Note

IVR-enabled VSANs can be configured when the interop mode is enabled (any interop mode) or disabled (no interop mode).

Domain ID Guidelines

Prior to Cisco MDS SAN-OS Release 2.1(1a), unique domain IDs are required across inter-connected VSANs. Consider the following guidelines for unique domain IDs:

- Minimize the number of switches that require a domain ID assignment. This ensures minimum traffic disruption.
- Minimize the coordination between interconnected VSANs, when configuring the SAN for the first time, as well as when you add each new switch.



Note

As of Cisco MDS SAN-OS Release 2.1(1a), unique domain IDs are no longer required.

Send documentation comments to mdsfeedback-doc@cisco.com.

Transit VSAN Guidelines

Consider the following guidelines for transit VSANs:

- Besides defining the IVZ membership, you can choose to specify a set of transit VSANs to provide connectivity between two edge VSANs:
 - If two edge VSANs in an IVZ overlap, then a transit VSAN is not required (though, not prohibited) to provide connectivity.
 - If two edge VSANs in an IVZ do not overlap, you may need one or more transit VSANs to provide connectivity. Two edge VSANs in an IVZ will not overlap if IVR is not enabled on a switch that is a member of both the source and destinations edge VSANs.
- Traffic between the edge VSANs only traverses through the shortest IVR path.
- Transit VSAN information is common to all IVZs. Sometimes, a transit VSAN can also double-up as an edge VSAN in another IVZ.

Border Switch Guidelines

Before configuring border switches, consider the following guidelines:

- Border switches require Cisco MDS SAN-OS Release 1.3(1) or later.
- A border switch must be a member of two or more VSANs.
- A border switch that facilitates IVR communications must be IVR enabled.
- IVR can (optionally) be enabled on additional border switches to provide redundant paths between active IVZ members.
- The VSAN topology configuration must be updated before a border switch is added or removed if the switch is running Cisco MDS SAN-OS Release 2.0(3) or earlier, or if IVR topology automatic mode is not enabled (available as of Cisco MDS SAN-OS Release 2.1(1a) or later).

IVR Configuration

To configure IVR in a SAN fabric, follow these steps:

-
- | | |
|---------------|--|
| Step 1 | Determine whether to use IVR NAT (Network Address Translation). |
| Step 2 | If you do not plan to use IVR NAT (supported as of Cisco MDS SAN-OS Release 2.1(1a)), verify that unique domain IDs are configured in all switches and VSANs participating in IVR. |
| Step 3 | Enable IVR in the border switches. |
| Step 4 | Configure the service group as required. |
| Step 5 | Configure fabric distribution as required. |
| Step 6 | Configure the IVR topology, either manually or automatically. |
| Step 7 | Create and activate IVZSs in <i>all</i> of the IVR-enabled border switches, either manually or using fabric distribution. |
| Step 8 | Verify the IVR configuration. |
-

Send documentation comments to mdsfeedback-doc@cisco.com.

Unique Domain ID Configuration Options

If you are not using IVR NAT, you must use unique domain IDs. You can configure unique domain IDs using one of two options:

- Configure the allowed-domains list so that the domains in different VSANs are non-overlapping on all participating switches and VSANs.
- Configure static, non-overlapping domains for each participating switch and VSAN (see [Chapter 31, “Configuring Domain Parameters”](#)).



Note

If you are using IVR NAT, you do not need unique domain IDs.

Enabling IVR

The IVR feature must be enabled in all border switches in the fabric that participate in the IVR. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family. Prior to Cisco MDS SAN-OS Release 2.0(1b), enabling IVR on all required switches in the fabric was performed manually. As of Cisco MDS SAN-OS Release 2.0(1b), you can configure fabric-wide distribution of the IVR configuration (“[IVR Configuration Distribution](#)” section on page 17-6).

The configuration and verification commands for the IVR feature are only available when IVR is enabled on a switch. When you disable this configuration, all related configurations are automatically discarded.

To enable IVR on any participating switch, follow these steps:

| | Command | Purpose |
|--------|--------------------------------------|---------------------------------------|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# ivr enable | Enables IVR on the switch. |
| | switch(config)# no ivr enable | Disables (default) IVR on the switch. |

IVR Configuration Distribution

The IVR feature uses the Cisco Fabric Services (CFS) infrastructure to enable efficient configuration management and to provide a single point of configuration for the entire fabric in the VSAN (see [Chapter 9, “Using the CFS Infrastructure”](#)).

The following configurations are distributed:

- IVR zones.
- IVR zone sets.
- IVR VSAN topology.
- IVR active topology and zone set (activating these features in one switch propagates the configuration to all other distribution-enabled switches in the fabric).
- IVR service groups.
- AFID database.

Send documentation comments to mdsfeedback-doc@cisco.com.

**Note**

IVR configuration distribution is disabled by default. For the feature to function correctly, you must enable it on all IVR-enabled switches in the network.

Database Implementation

The IVR feature uses three databases to accept and implement configurations.

- Configured database—The database is manually configured by the user.
- Active database—The database is currently enforced by the fabric.
- Pending database—If you modify the configuration, you need to commit or discard the configured database changes to the pending database. The fabric remains locked during this period. Changes to the pending database are not reflected in the active database until you commit the changes to CFS.

Enabling Configuration Distribution

To enable IVR configuration distribution, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ivr distribute switch(config)# no ivr distribute | Enables IVR distribution. Disables (default) IVR distribution. |

Locking the Fabric

The first action that modifies the database creates the pending database and locks the feature in the VSAN. Once you lock the fabric, the following situations apply:

- No other user can make any configuration changes to this feature.
- A copy of the configuration database becomes the pending database along with the first active change.

Committing the Changes

If you commit the changes made to the active database, the configuration is committed to all the switches in the fabric. On a successful commit, the configuration change is applied throughout the fabric and the lock is released.

To commit IVR configuration changes, follow these steps:

| | Command | Purpose |
|--------|--|----------------------------|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ivr commit | Commits the IVR changes. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Discarding the Changes

If you discard (abort) the changes made to the pending database, the configuration database remains unaffected and the lock is released.

To discard IVR configuration changes, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ivr abort | Discards the IVR changes and clears the pending configuration database. |

Clearing a Locked Session

If you have performed an IVR task and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.



Tip

The pending database is only available in the volatile directory and are subject to being discarded if the switch is restarted.

To use administrative privileges and release a locked DPVM session, use the **clear ivr session** command in EXEC mode.

```
switch# clear ivr session
```

About IVR NAT

Prior to Cisco MDS SAN-OS Release 2.1(1a), IVR required unique domain IDs for all switches in the fabric. As of Cisco MDS SAN-OS Release 2.1(1a), you can enable IVR Network Address Translation (NAT) to allow non-unique domain IDs. This feature simplifies the deployment of IVR in an existing fabric where non-unique domain IDs might be present.

IVR NAT is enabled in all border IVR-enabled switches in the fabric IVR configuration distribution (see the “[IVR Configuration Distribution](#)” section on page 17-6). By default, IVR NAT, and IVR configuration distribution are disabled in all switches in the Cisco MDS 9000 Family.



Note

For IVR NAT to function correctly in the network, all IVR-enabled switches must run Cisco MDS SAN-OS Release 2.1(1a) or later.

IVR NAT virtualizes the switches in other VSANs by using local VSAN for the destination IDs in the Fibre Channel headers. In some message types, the destinations IDs are part of the payload. In these cases, IVR NAT replaces the actual destination ID with the virtualized destination ID. IVR NAT supports destination ID replacement in the messages described in [Table 17-1](#).

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 17-1 **Extended Link Service Messages Supported by IVR NAT**

| Extended Link Service Messages | Link Service Command (LS_COMMAND) | Mnemonic |
|---|-----------------------------------|------------|
| Abort Exchange | 0x06 00 00 00 | ABTX |
| Discover Address | 0x52 00 00 00 | ADISC |
| Discover Address Accept | 0x02 00 00 00 | ADISC ACC |
| Fibre Channel Address Resolution Protocol Reply | 0x55 00 00 00 | FARP-REPLY |
| Fibre Channel Address Resolution Protocol Request | 0x54 00 00 00 | FARP-REQ |
| Logout | 0x05 00 00 00 | LOGO |
| Port Login | 0x30 00 00 00 | PLOGI |
| Read Exchange Concise | 0x13 00 00 00 | REC |
| Read Exchange Concise Accept | 0x02 00 00 00 | REC ACC |
| Read Exchange Status Block | 0x08 00 00 00 | RES |
| Read Exchange Status Block Accept | 0x02 00 00 00 | RES ACC |
| Read Link Error Status Block | 0x0F 00 00 00 | RLS |
| Read Sequence Status Block | 0x09 00 00 00 | RSS |
| Reinstate Recovery Qualifier | 0x12 00 00 00 | RRQ |
| Request Sequence Initiative | 0x0A 00 00 00 | RSI |
| Scan Remote Loop | 0x7B 00 00 00 | RSL |
| Third Party Process Logout | 0x24 00 00 00 | TPRLO |
| Third Party Process Logout Accept | 0x02 00 00 00 | TPRLO ACC |

If you have a message that is not recognized by IVR NAT and contains the destination ID in the payload, you cannot use IVR with NAT in your topology. You can still use IVR with unique domain IDs.

Enabling IVR NAT

The configuration and verification commands for the IVR feature are only available when IVR is enabled on a switch. When you disable this configuration, all related configurations are automatically discarded.

To configure IVR NAT, follow these steps:

| | Command | Purpose |
|---------------|--|---|
| Step 1 | switch# confi g t | Enters configuration mode. |
| Step 2 | switch(config)# ivr fcid-nat | Enables IVR NAT on the switch. |
| | switch(config)# no ivr fcid-nat | Disables (default) IVR NAT on the switch. |

Send documentation comments to mdsfeedback-doc@cisco.com.

About IVR Topologies

IVR must know about the topology of the IVR-enabled switches in the fabric to function properly. You can specify the topology two ways:

- Manual configuration

If you manually configure the IVR topology, you must ensure that the IVR topology exists on every IVR-enabled switch in the fabric. You can configure the IVR topology manually on each IVR-enabled switch or, as of Cisco MDS SAN-OS Release 2.0(1b), you can use CFS to distribute the configuration automatically (see the “[Database Merge Guidelines](#)” section on page 17-23).

If an IVR-enabled switch is removed from the network, the IVR topology database must be updated to reflect the change.

- Automatic mode

As of Cisco MDS SAN-OS Release 2.1(1a)), you can configure IVR topology automatic mode. Automatic mode uses CFS configuration distribution to dynamically learn and maintain up-to-date information about the topology of the IVR-enabled switches in the network.

If a manually configured IVR topology database exists, automatic mode initially uses that topology information. This reduces disruption in the network by gradually migrating from a user-specified topology database to an automatically learned topology database. Then the user-configured topology entries that are not part of the network are aged out in about three minutes and new entries that are not part of user configured database are added as they are learned from the network.

Configuring IVR Topologies

This section describes how to manually configure an IVR topology or how to configure IVR topology automatic mode.

Manually Configuring the IVR Topology

You can have up to 64 VSANs (or 128 VSANs as of Cisco MDS SAN-OS Release 2.1(1a)) in an IVR topology. Specify the IVR topology using the following information:

- The switch WWNs of the IVR-enabled switches.
- A minimum of two VSANs to which the IVR-enabled switch belongs.
- The autonomous fabric ID (AFID), which distinguishes two VSANs that are logically and physically separate, but have the same VSAN number. Cisco MDS SAN-OS Releases 1.3(1) through 2.0(2b) support only one default AFID (AFID 1) and thus does not support non-unique VSAN IDs in the network. As of Cisco MDS SAN-OS Release 2.1(1a), you can specify up to 64 AFIDs. See [Figure 17-2](#).




Note

If two VSANs in an IVR topology have the same VSAN ID and different AFIDs, they count as two VSANs for the 128-VSAN limit for IVR.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 17-2 Example IVR Topology with Non-Unique VSAN IDs Using AFIDs



Note

The use of a single AFID does not allow for VSANs that are logically and physically separate but have the same VSAN number in an IVR topology.



Caution

You can only configure a maximum of 128 IVR-enabled switches and 64 distinct VSANs (or 128 distinct VSANs as of Cisco MDS SAN-OS Release 2.1(1a)) in an IVR topology (see the [“Database Merge Guidelines”](#) section on page 17-23).

Configuring an IVR Topology Database

Use the **show wwn switch** command to obtain the switch WWNs of the IVR-enabled switches.

To configure a user-defined IVR topology database, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# ivr vsan-topology database switch(config-ivr-topology-db)# | Enters the VSAN topology database configuration mode for the IVR feature. |
| Step 3 | switch(config-ivr-topology-db)# autonomous-fabric-id 1 switch 20:00:00:05:30:01:1b:b8 vsan-ranges 1-2,6 | Configures VSANs 1, 2, and 6 to participate in IVR for this switch. |
| Step 4 | switch(config-ivr-topology-db)# autonomous-fabric-id 1 switch 20:00:00:05:30:01:1b:c2 vsan-ranges 1-3 | Configures VSANs 1, 2 and 3 to participate in IVR for this switch. |
| Step 5 | switch(config-ivr-topology-db)# no autonomous-fabric-id 1 switch 20:00:00:05:30:01:1b:c2 vsan-ranges 1-2 | Removes VSANs 1 and 2 from IVR for this switch. |
| Step 6 | switch(config-ivr-topology-db)# end switch# | Reverts to EXEC mode. |

Send documentation comments to mdsfeedback-doc@cisco.com.

View your configured IVR topology using the **show ivr vsan-topology** command. In the following example output, VSAN 2 is the transit VSAN between VSANs 1, 5, and 6.

```
switch# show ivr vsan-topology
```

| AFID | SWITCH WWN | Active | Cfg. VSANS |
|------|---------------------------|--------|------------|
| 1 | 20:00:00:05:30:01:1b:c2 * | no | yes 1-2 |
| 1 | 20:02:00:44:22:00:4a:05 | no | yes 1-2,6 |
| 1 | 20:02:00:44:22:00:4a:07 | no | yes 2-5 |

Total: 3 entries in active and configured IVR VSAN-Topology

Current Status: Inter-VSAN topology is INACTIVE



Note

If CFS is not enabled, you must repeat this configuration in all IVR-enabled switches. See the [“Database Merge Guidelines”](#) section on page 17-23.



Tip

Transit VSANs are deduced based on your configuration. The IVR feature does not have an explicit transit VSAN configuration.

Activating a Manually Configured IVR Topology

After manually configuring the IVR topology database, you must activate it.



Caution

Active IVR topologies cannot be deactivated. You can only switch to IVR topology automatic mode.

To activate a manually configured IVR topology, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ivr vsan-topology activate | Activates the configured IVR topology. |

View your active IVR topology using the **show ivr vsan-topology** command.

```
switch# show ivr vsan-topology
```

| AFID | SWITCH WWN | Active | Cfg. VSANS |
|------|---------------------------|--------|------------|
| 1 | 20:00:00:05:30:01:1b:c2 * | yes | yes 1-2 |
| 1 | 20:02:00:44:22:00:4a:05 | yes | yes 1-2,6 |
| 1 | 20:02:00:44:22:00:4a:07 | yes | yes 2-5 |

Total: 3 entries in active and configured IVR VSAN-Topology

Current Status: Inter-VSAN topology is **ACTIVE**

Last activation time: Mon Mar 24 07:19:53 1980

Send documentation comments to mdsfeedback-doc@cisco.com.

Configuring IVR Topology Automatic Mode

As of Cisco MDS SAN-OS Release 2.1(1a), you can configure IVR topology automatic mode.



Note

IVR configuration distribution must be enabled before configuring IVR topology automatic mode (see the “[IVR Configuration Distribution](#)” section on page 17-6). Once IVR topology automatic mode is enabled, you cannot disable IVR configuration distribution.

To configure IVR topology automatic mode, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ivr vsan-topology auto | Configures IVR topology automatic mode. |
| | switch(config)# ivr vsan-topology activate | Disables IVR topology automatic mode and reverts to user-configuration mode. |

View automatically discovered IVR topology using the **show ivr vsan-topology** command.

```
switch# show ivr vsan-topology
AFID  SWITCH WNN                Active  Cfg. VSANS
-----
1  20:00:00:05:30:01:1b:c2 *  yes    yes    1-2
1  20:02:00:44:22:00:4a:05    yes    yes    1-2,6
1  20:02:00:44:22:00:4a:07    yes    yes    2-5

Total:   3 entries in active and configured IVR VSAN-Topology

Current Status: Inter-VSAN topology is AUTO
Last activation time: Mon Mar 24 07:19:53 1980
```

Migrating from IVR Topology Automatic Mode to Manual Mode

If you want to migrate the active IVR VSAN topology database from automatic mode to user-configured mode, first copy the active IVR VSAN topology database to the user-configured IVR VSAN topology database before switching modes. To migrate from automatic mode to manual mode, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# ivr copy auto-topology user-configured-topology | Copies the automatic IVR topology database to the user-configured IVR topology. |
| Step 2 | switch# config t switch(config)# | Enters configuration mode. |
| Step 3 | switch(config)# ivr vsan-topology active | Disabled automatic mode for the IVR topology database and enables user-configuration mode. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Clearing the Configured IVR Topology Database

To clear the user-configured IVR VSAN topology database using, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# no ivr vsan-topology database | Clears the previously created IVR topology. |

Non-Unique VSAN IDs Using AFIDs

As of Cisco MDS SAN-OS Release 2.1(1a), you can configure more than one AFID. This feature allows more than one VSAN in the network with the same VSAN ID. Using this feature you can avoid downtime when enabling IVR between fabrics that contain VSANs with the same ID. However, for VSANs with the same ID to communicate, there must be a transit VSAN with a different VSAN ID between the source and target VSANs.



Note

AFID configuration is used only when the VSAN topology mode is automatic. In user-configured VSAN topology mode, the AFIDs are specified in the VSAN topology configuration itself and a separate AFID configuration is not needed.

Configuring the AFID Database

To configure the AFID database, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# autonomous-fabric-id database | Enters AFID database configuration submenu. |
| Step 3 | switch(config-afid-db)# switch-wwn 20:00:00:0c:91:90:3e:80 autonomous-fabric-id 10 vsan-ranges 1,2,5-8 | Configures an AFID and VSAN range for a switch. |
| | switch(config-afid-db)# no switch-wwn 20:00:00:0c:91:90:3e:80 autonomous-fabric-id 2 vsan-ranges 3 | Deletes VSAN 3 from AFID 2. |
| Step 4 | switch(config-afid-db)# switch-wwn 20:00:00:0c:91:90:3e:80 default-autonomous-fabric-id 5 | Configures the default AFID for all VSANs not explicitly associated with an AFID. |
| Step 5 | switch(config-afid-db)# no switch-wwn 20:00:00:0c:91:90:3e:80 default-autonomous-fabric-id 5 | Deletes the default AFID. |

Send documentation comments to mdsfeedback-doc@cisco.com.

View the contents of the AFID database using the **show autonomous-fabric-id database** command.

```
switch# show autonomous-fabric-id database
```

```
SWITCH WWN                                Default-AFID
-----
20:00:00:0c:91:90:3e:80                    5
```

Total: 1 entry in default AFID table

```
SWITCH WWN                                AFID      VSANS
-----
20:00:00:0c:91:90:3e:80                    10      1,2,5-8
```

Total: 1 entry in AFID table

Adding IVR Virtual Domain

In a remote VSAN, the IVR application does not automatically add the virtual domain to the assigned domain list. Some switches (for example, the Cisco SN5428) do not query the remote name server until the remote domain appears in the assigned domain list in the fabric. In such cases, add the IVR virtual domains in a specific VSAN(s) to the assigned domain list in that VSAN. When adding IVR domains, all IVR virtual domains that are currently present in the fabric (and any virtual domain that is created in the future) will appear in the assigned domain list for that VSAN.



Tip

Be sure to add IVR virtual domains if Cisco SN5428 or Cisco MDS 9020 switches exist in the VSAN.



Tip

As of Cisco MDS SAN-OS Release 1.3(4a), only add IVR domains in the edge VSANs and not in transit VSANs.

To add an IVR virtual domain to a specified VSAN, follow these steps:

| | Command | Purpose |
|---------------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ivr virtual-fcdomain-add vsan-ranges 1 | Adds the IVR virtual domains in VSAN 1. |
| Step 3 | switch(config)# no ivr virtual-fcdomain-add vsan-ranges 1 | Reverts to the factory default of not adding IVR virtual domains and removes the currently active virtual domains for that VSAN from the fcdomain manger list |

View the status of the IVR virtual domain configuration using the **show ivr virtual-fcdomain-add-status** command.

```
switch# show ivr virtual-fcdomain-add-status
IVR virtual domains are added to fcdomain list in VSANS: 1
(As well as to VSANS in interoperability mode 2 or 3)
```

Send documentation comments to mdsfeedback-doc@cisco.com.

When you enable the IVR virtual domains, links may fail to come up due to overlapping virtual domain identifiers. If so, temporarily withdraw the overlapping virtual domain from that VSAN.



Note

Withdrawing an overlapping virtual domain from an IVR VSAN disrupts IVR traffic to and from that domain.

Use the **ivr withdraw domain** command in EXEC mode to temporarily withdraw the overlapping virtual domain interfaces from the affected VSAN.

About IVZs and IVZSs

As part of the IVR configuration, you need to configure one or more IVZs to enable cross-VSAN communication. To achieve this result, you must specify each IVZ as a set of (pWWN, VSAN) entries. Like zones, several IVZs can be configured to belong to an IVR zone. You can define several IVZSs and activate only one of the defined IVZSs.



Note

The same IVZS must be activated on *all* of the IVR-enabled switches.



Caution

As of Cisco MDS SAN-OS Release 1.3(1b), you can only configure a total number of 2000 zone members on all switches in a network. As of Cisco MDS SAN-OS Release 2.1(1a), the limit is increased to a total number of 10,000 zone members on all switches in a network. A zone member is counted twice if it exists in two zones. See the [“Database Merge Guidelines”](#) section on page 17-23.

IVZs Versus Zones

Table 17-2 identifies the key differences between IVZs and zones.

Table 17-2 Key Differences between IVZs and Zones

| IVZs | Zones |
|--|--|
| IVZ membership is specified using the VSAN and pWWN combination. | Zone membership is specified using pWWN, fabric WWN, sWWN, or the fabric ID. |
| Default zone policy is always deny (not configurable). | Default zone policy is deny (configurable). |

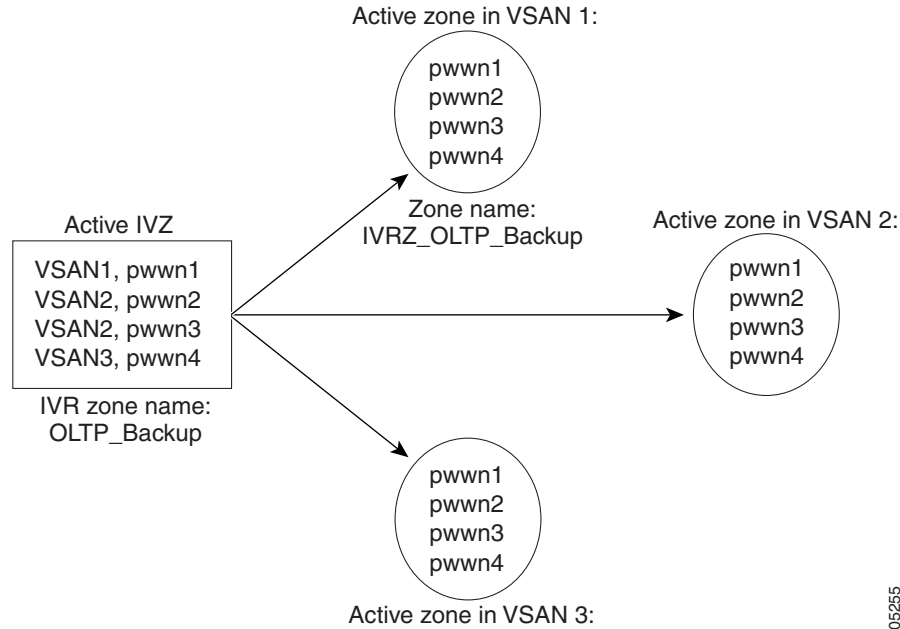
Automatic IVZ Creation

Figure 17-3 depicts an IVZ consisting of four members. To allow pwnn1 to communicate with pwnn2, they must be in the same zone in VSAN 1, as well as in VSAN 2. If they are not in the same zone, then the hard-zoning ACL entries will prohibit pwnn1 from communicating with pwnn2.

A zone corresponding to each active IVZ is automatically created in each edge VSAN specified in the active IVZ. All pWWNs in the IVZ are members of these zones in each VSAN.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 17-3 Creating Zones on IVZ Activation



The zones are created automatically by the IVR process when an IVZS is activated. They are not stored in full zone set database and are lost when the switch reboots or when a new zone set is activated. The IVR feature monitors these events and adds the zones corresponding to the active IVZS configuration when a new zone set is activated. Like zone sets, IVR zone sets are also activated nondisruptively.



Note

If pwwn1 and pwwn2 are in an IVZ in the current as well as the new IVZS, then activation of the new IVZS does not cause any traffic disruption between them.

IVZ and IVZS names are restricted to 64 alphanumeric characters.



Caution

You can only configure a total of 200 zones and 32 zone sets on the switches in the network. As of Cisco MDS SAN Release 2.1(1a), you can configure up to 2000 zones on the switches in the network. See the [“Database Merge Guidelines”](#) section on page 17-23.

[Send documentation comments to mdsfeedback-doc@cisco.com.](mailto:mdsfeedback-doc@cisco.com)

Configuring IVZs and IVZSs

This section describes how to configure IVZs and IVZSs.

Creating and Activating IVZs and IVZSs

To create and activate IVZs and IVZSs, follow these steps:

| | Command | Purpose |
|---------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# ivr zone name sample_vsan2-3 switch(config-ivr-zone)# | Creates an IVR zone named sample_vsan2-3. |
| Step 3 | switch(config-ivr-zone)# member pwwn 21:00:00:e0:8b:02:ca:4a vsan 3 | Adds the specified pWWN in VSAN 3 as an IVZ member. |
| Step 4 | switch(config-ivr-zone)# member pwwn 21:00:00:20:37:c8:5c:6b vsan 2 | Adds the specified pWWN in VSAN 2 as an IVZ member. |
| Step 5 | switch(config-ivr-zone)# exit switch(config)# | Reverts to configuration mode. |
| Step 6 | switch(config)# ivr zone name sample_vsan4-5 switch(config-ivr-zone)# | Creates an IVZ named sample_vsan4-5. |
| Step 7 | switch(config-ivr-zone)# member pwwn 21:00:00:e0:8b:06:d9:1d vsan 4 | Adds the specified pWWN in VSAN 4 as an IVZ member. |
| Step 8 | switch(config-ivr-zone)# member pwwn 21:01:00:e0:8b:2e:80:93 vsan 4 | Adds the specified pWWN in VSAN 4 as an IVZ member. |
| Step 9 | switch(config-ivr-zone)# member pwwn 10:00:00:00:c9:2d:5a:dd vsan 5 | Adds the specified pWWN in VSAN 5 as an IVZ member. |
| Step 10 | switch(config-ivr-zone)# exit switch(config)# | Reverts to configuration mode. |
| Step 11 | switch(config)# ivr zoneset name Ivr_zoneset1 switch(config-ivr-zoneset)# | Creates an IVZS named Ivr_zoneset1. |
| Step 12 | switch(config-ivr-zoneset)# member sample_vsan2-3 | Adds the sample_vsan2-3 IVZ as an IVZS member. |
| Step 13 | switch(config-ivr-zoneset)# member sample_vsan4-5 | Adds the sample_vsan4-5 IVZ as an IVZS member. |
| Step 14 | switch(config-ivr-zoneset)# exit switch(config)# | Returns to configuration mode. |
| Step 15 | switch(config)# ivr zoneset activate name IVR_ZoneSet1 | Activates the newly created IVZS. |
| | switch(config)# ivr zoneset activate name IVR_ZoneSet1 force | Forcefully activates the specified IVZS. |
| | switch(config)# no ivr zoneset activate name IVR_ZoneSet1 | Deactivates the specified IVZS. |
| Step 16 | switch(config-ivr-zoneset)# end switch# | Returns to EXEC mode. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Configuring LUNs in IVR Zoning

LUN zoning can be used between members of active IVZs. Prior to Cisco MDS SAN-OS Release 2.1(1a), you can configure the service by creating and activating LUN zones between the desired IVZ members in all relevant edge VSANs using the zoning interface. As of Cisco MDS SAN-OS Release 2.1(1a), IVR directly supports LUN zoning. For more details on the advantages of LUN zoning, see the [“About LUN Zoning” section on page 15-17](#).

To configure LUNs in IVR zoning in Cisco MDS SAN-OS Release 2.1(1a) or later, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ivr zone name IvrLunZone switch(config-ivr-zone)# | Configures an IVZ called IvrLunZone. |
| Step 3 | switch(config-ivr-zone)# member pwwn 10:00:00:23:45:67:89:ab lun 64 vsan 10 | Configures an IVZ member based on the specified pWWN and LUN value. |
| | switch(config-ivr-zone)# member pwwn 10:00:00:23:45:67:89:ab lun 64 vsan 10 autonomous-fabric-id 20 | Configures an IVZ member based on the specified pWWN, LUN value, and AFID. |
| | switch(config-ivr-zone)# no member pwwn 20:81:00:0c:85:90:3e:80 lun 32 vsan 13 autonomous-fabric-id 10 | Removes an IVZ member. |



Note

As of Cisco MDS SAN-OS Release 2.1(1a), you can configure LUN zoning in an IVZS setup.

Configuring the QoS Attribute

As of Cisco MDS SAN-OS Release 2.1(1a), you can configure a QoS attribute for an IVZ. To configure QoS for an IVSZ, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ivr zone name IvrZone switch(config-ivr-zone)# | Configures an IVZ called IvrZone. |
| Step 3 | switch(config-ivr-zone)# attribute qos priority medium | Configures the QoS for IVZ traffic to medium. |
| | switch(config-ivr-zone)# no attribute qos priority medium | Reverts to the default QoS setting. The default is low. |



Note

If other QoS attributes are configured, the highest setting takes priority.

Send documentation comments to mdsfeedback-doc@cisco.com.

Using the force Option

Use the **force** option to activate the specified IVZS. Table 17-3 lists the various scenarios with and without the **force** option.

Table 17-3 *IVR Scenarios with and without the force Option.*

| Case | Default Zone Policy | Active Zone Set before IVR Zone Activation | force Option Used? | IVZS Activation Status | Active IVR Zone Created? | Possible Traffic Disruption |
|----------------|---------------------|---|--------------------|------------------------|--------------------------|-----------------------------|
| 1 | Deny | No active zone set | No | Failure | No | No |
| 2 | | | Yes | Success | Yes | No |
| 3 ¹ | Deny | Active zone set present | No/Yes | Success | Yes | No |
| 4 | Permit | No active zone set or Active zone set present | No | Failure | No | No |
| 5 | | | Yes | Success | Yes | Yes |

1. We recommend that you use the Case 3 scenario.



Caution

Using the **force** option of IVZS activation may cause traffic disruption, even for devices that are not involved in IVR. For example, if your configuration does not have any active zone sets and the default zone policy is `permit`, then an IVZS activation will fail. However, IVZS activation will go through if the **force** option is used. Because zones are created in the edge VSANs corresponding to each IVZ, traffic may be disrupted in edge VSANs where the default zone policy is `permit`.

Clearing the IVZ Database



Note

Clearing a zone set only erases the configured zone database, not the active zone database.

To clear the IVZ database, use the **clear ivr zone database** command.

```
switch# clear ivr zone database
```

This command clears all configured IVZ information.



Note

After issuing a **clear ivr zone database** command, you need to explicitly issue the **copy running-config startup-config** to ensure that the running configuration is used when you next start the switch.

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying IVZ Configurations

View your IVZ configuration using the **show ivr zone** command.

```
switch# show ivr zone

zone name sample_vsan2-3
  pwwn 21:00:00:e0:8b:02:ca:4a vsan 3
  pwwn 21:00:00:20:37:c8:5c:6b vsan 2

zone name sample_vsan4-5
  pwwn 21:00:00:e0:8b:06:d9:1d vsan 4
  pwwn 21:01:00:e0:8b:2e:80:93 vsan 4
  pwwn 10:00:00:00:c9:2d:5a:dd vsan 5
```

View your IVZS configuration using the **show ivr zoneset** command.

```
switch# show ivr zoneset

zoneset name ivr_qa_zs_all
  zone name sample_vsan2-3
    pwwn 21:00:00:e0:8b:02:ca:4a vsan 3
    pwwn 21:00:00:20:37:c8:5c:6b vsan 2

  zone name sample_vsan4-5
    pwwn 21:00:00:e0:8b:06:d9:1d vsan 4
    pwwn 21:01:00:e0:8b:2e:80:93 vsan 4
    pwwn 10:00:00:00:c9:2d:5a:dd vsan 5
```

Use the **show ivr zoneset active** command to view your active IVZS status.

```
switch# show ivr zoneset active

zoneset name ivr_qa_zs_all
  zone name sample_vsan2-3
    * pwwn 21:00:00:e0:8b:02:ca:4a vsan 3
    pwwn 21:00:00:20:37:c8:5c:6b vsan 2

  zone name sample_vsan4-5
    pwwn 21:00:00:e0:8b:06:d9:1d vsan 4
    * pwwn 21:01:00:e0:8b:2e:80:93 vsan 4
    pwwn 10:00:00:00:c9:2d:5a:dd vsan 5
```



Tip

Repeat this configuration in all border switches participating in the IVR configuration.



Note

Using the Cisco MDS Fabric Manager, you can distribute IVZ configurations to all IVR-capable switches in the interconnected VSAN network. Refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*.

About IVR Service Groups

In a complex network topology, you might only have a few IVR-enabled VSANs. To reduce the amount of traffic to non-IVR-enabled VSANs, you can configure a service group that restricts the traffic to the IVR-enabled VSANs. Only one service group is allowed in a network. When a new IVR-enabled switch is added to the network, you must update the service group to include the new VSANs.

Send documentation comments to mdsfeedback-doc@cisco.com.


Note

CFS distribution of IVR information is restricted within the service group only when IVR VSAN topology is in automatic mode. See the [“About IVR Topologies”](#) section on page 17-10.

Configuring IVR Service Groups

To configure an IVR service group, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ivr service-group name IVR-SG1 switch(config-ivr-sg)# | Configures the IVR service group called IVR-SG1. |
| | switch(config)# no ivr service-group name IVR-SG1 | Deletes the IVR service group. |
| Step 3 | switch(config-ivr-sg)# autonomous-fabric-id 10 vsan-ranges 1,2,6-10 | Configures AFID 10 for VSANs 1, 2, and 6 through 10. |
| | switch(config-ivr-sg)# no autonomous-fabric-id 20 vsan-ranges 1,2,6-10 | Removes the association between AFID 20 and VSANs 1, 2, and 6 through 10. |

Use the **show ivr service-group database** command to view the IVR service group database configuration.

```
switch(config)# show ivr service-group database
```

| SG-ID | SG-NAME | AFID | VSANS |
|-------|---------|------|-----------|
| 1 | IVR-SG1 | 10 | 1-2, 6-10 |
| 1 | IVR-SG1 | 11 | 1 |

Total: 2 entries in service group table

IVR Interoperability

When using the IVR feature, all border switches in a given fabric must be Cisco MDS switches. However, other switches in the fabric may be non-MDS switches. For example, end devices that are members of the active IVZS may be connected to non-MDS switches. Non-MDS switches may also be present in the transit VSAN(s) or in the edge VSANs if one of the **interop** modes is enabled.

See the [“Switch Interoperability”](#) section on page 39-22.

Configuring IVR Using Read-Only Zoning

Read-only zoning (with or without LUNs) can be used between members of active IVR zones. To configure this service, you must create and activate read-only zones between the desired IVZ members in all relevant edge VSANs using the zoning interface.

Send documentation comments to mdsfeedback-doc@cisco.com.



Note

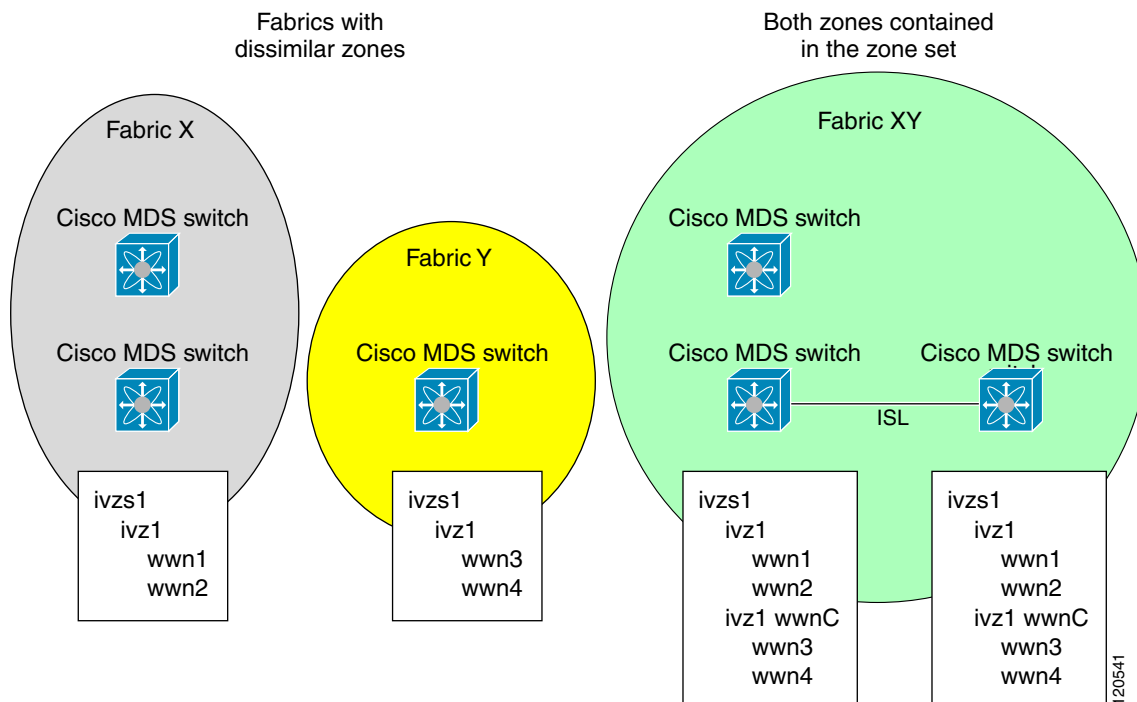
Read-only zoning cannot be configured in an IVZS setup.

Database Merge Guidelines

A database merge refers to a union of the configuration database and static (unlearned) entries in the active database. Refer to the “[CFS Merge Support](#)” section on page 9-6 for detailed concepts.

- Be aware of the following conditions when merging two IVR fabrics:
 - The IVR configurations are merged even if two fabrics contain different configurations.
 - If dissimilar zones exist in two merged fabrics, the zone from each fabric is cloned in the distributed zone set with appropriate names (see [Figure 17-4](#)).

Figure 17-4 Fabric Merge Consequences



- You can configure different IVR configurations in different Cisco MDS switches.
- Be aware that the merge follows more liberal approach in order to avoid traffic disruption. After the merge, the configuration will be a union of the configurations that were present on the two switches involved in the merge.
 - The configurations are merged even if both fabrics have different configurations.
 - A union of zones and zone sets are used to get the merged zones and zone sets. If a dissimilar zone exists in two fabrics, the dissimilar zones are cloned into the zone set with appropriate names so both zones are present.
 - The merged topology contains a union of the topology entries for both fabrics.

Send documentation comments to mdsfeedback-doc@cisco.com.

- The merge will fail if the merged database contains more topology entries than the allowed maximum.
- The total number of VSANs across the two fabrics cannot exceed 64. As of Cisco MDS SAN-OS Release 2.1(1a), the total number of VSANs across the two fabrics cannot exceed 128.



Note VSANs with the same VSAN ID but different AFIDs are counted as two separate VSANs.

- The total number of IVR-enabled switches across the two fabrics cannot exceed 128.
- The total number of zone members across the two fabrics cannot exceed 2000. As of Cisco MDS SAN-OS Release 2.1(1a), the total number of zone members across the two fabrics cannot exceed 10,000. A zone member is counted twice if it exists in two zones.
- The total number of zones across the two fabrics cannot exceed 200. As of Cisco MDS SAN-OS Release 2.1(1a), the total number of zones across the two fabrics cannot exceed 2000.
- The total number of zone sets across the two fabrics cannot exceed 32.

Table 17-4 describes the results of a CFS merge of two IVR-enabled fabrics under different conditions.

Table 17-4 Results of Merging Two IVR-Enabled Fabrics

| IVR Fabric 1 | IVR Fabric 2 | After Merge |
|--|-----------------|--|
| NAT enabled | NAT disabled | Merge succeeds and NAT enabled |
| Auto mode on | Auto mode off | Merge succeeds and auto mode on |
| Conflicting AFID database | | Merge fails |
| Conflicting IVR zone set database | | Merge succeeds with new zones created to resolve conflicts |
| Combined configuration exceeds limits (such as maximum number of zones or VSANs) | | Merge fails |
| Service group 1 | Service group 2 | Merge succeeds with service groups combined |
| User-configured VSAN topology configuration with conflicts | | Merge fails |
| User-configured VSAN topology configuration without conflicts | | Merge succeeds |



Caution

If you do not follow these conditions, the merge will fail. The next distribution will forcefully synchronize the databases and the activation states in the fabric.

Send documentation comments to mdsfeedback-doc@cisco.com.

Configuring IVR Logging Levels

To configure the severity level for logging messages from the IVR feature, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# logging level ivr 4 | Configures Telnet or SSH logging for the IVR feature at level 4 (warning). As a result, logging messages with a severity level of 4 or above are displayed. |

Use the **show logging level** command to view the configured logging level for the IVR feature.

```
switch# show logging level
Facility           Default Severity      Current Session Severity
-----
...
ivr                5                      4
...
0(emergencies)      1(alerts)              2(critical)
3(errors)            4(warnings)            5(notifications)
6(information)      7(debugging)
```

Displaying IVR Information

You can verify IVR information by using the **show ivr** set of commands. If you request information for a specific object (for example, a specific zone, zone set, VSAN, alias, or even a keyword like **brief** or **active**), only information for the specified object is displayed. If you do not request specific information, all available information is displayed. See Examples 17-1 to 17-12.

Example 17-1 Displays the Configured IVR VSAN Topology

```
switch# show ivr vsan-topology
AFID    SWITCH WWN           Active   Cfg. VSANS
-----
120:00:00:05:30:01:1b:c2 * yes      yes 1-2
1  20:02:00:44:22:00:4a:05 yes      yes 1-2,6
1  20:02:00:44:22:00:4a:07 yes      yes 2-5

Total:   5 entries in active and configured IVR VSAN-Topology

Current Status: Inter-VSAN topology is ACTIVE
Last activation time: Sat Mar 22 21:46:15 1980
```

The asterisk (*) indicates the local switch.

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 17-2 Displays the Active IVR VSAN Topology

```
switch# show ivr vsan-topology active
AFID  SWITCH WNN                Active   Cfg.  VSANS
-----
    1  20:00:00:05:30:01:1b:c2 *   yes     yes   1-2
    1  20:02:00:44:22:00:4a:05     yes     yes   1-2,6
    1  20:02:00:44:22:00:4a:07     yes     yes   2-5

Total:    5 entries in active IVR VSAN-Topology

Current Status: Inter-VSAN topology is ACTIVE
Last activation time: Sat Mar 22 21:46:15
```

Example 17-3 Displays the Configured IVR VSAN Topology

```
switch# show ivr vsan-topology configured
AFID  SWITCH WNN                Active   Cfg.  VSANS
-----
    1  20:00:00:05:30:01:1b:c2 *   yes     yes   1-2
    1  20:02:00:44:22:00:4a:05     yes     yes   1-2,6
    1  20:02:00:44:22:00:4a:07     yes     yes   2-5

Total:    5 entries in configured IVR VSAN-Topology
```

Example 17-4 Displays the IVZ Configuration

```
switch# show ivr zone
zone name sample_vsan2-3
  pwwn 21:00:00:e0:8b:02:ca:4a vsan 3
  pwwn 21:00:00:20:37:c8:5c:6b vsan 2

zone name ivr_qa_z_all
  pwwn 21:00:00:e0:8b:06:d9:1d vsan 1
  pwwn 21:01:00:e0:8b:2e:80:93 vsan 4
  pwwn 10:00:00:00:c9:2d:5a:dd vsan 1
  pwwn 10:00:00:00:c9:2d:5a:de vsan 2
  pwwn 21:00:00:20:37:5b:ce:af vsan 6
  pwwn 21:00:00:20:37:39:6b:dd vsan 6
  pwwn 22:00:00:20:37:39:6b:dd vsan 3
  pwwn 22:00:00:20:37:5b:ce:af vsan 3
  pwwn 50:06:04:82:bc:01:c3:84 vsan 5
```

Example 17-5 Displays the Active IVZS Configuration

```
switch# show ivr zoneset active
zoneset name IVR_ZoneSet1
  zone name sample_vsan2-3
    pwwn 21:00:00:e0:8b:02:ca:4a vsan 3
    pwwn 21:00:00:20:37:c8:5c:6b vsan 2
```

Example 17-6 Displays Information for a Specified IVZ

```
switch# show ivr zone name sample_vsan2-3
zone name sample_vsan2-3
  pwwn 21:00:00:e0:8b:02:ca:4a vsan 3
  pwwn 21:00:00:20:37:c8:5c:6b vsan 2
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 17-7 Displays the Specified Zone in the Active IVZS

```
switch# show ivr zone name sample_vsan2-3 active
zone name sample_vsan2-3
  pwwn 21:00:00:e0:8b:02:ca:4a vsan 3
  pwwn 21:00:00:20:37:c8:5c:6b vsan 2
```

Example 17-8 Displays the IVZS Configuration

```
switch# show ivr zoneset
zoneset name ivr_qa_zs_all
  zone name ivr_qa_z_all
    pwwn 21:00:00:e0:8b:06:d9:1d vsan 1
    pwwn 21:01:00:e0:8b:2e:80:93 vsan 4
    pwwn 10:00:00:00:c9:2d:5a:dd vsan 1
    pwwn 10:00:00:00:c9:2d:5a:de vsan 2
    pwwn 21:00:00:20:37:5b:ce:af vsan 6
    pwwn 21:00:00:20:37:39:6b:dd vsan 6
    pwwn 22:00:00:20:37:39:6b:dd vsan 3
    pwwn 22:00:00:20:37:5b:ce:af vsan 3
    pwwn 50:06:04:82:bc:01:c3:84 vsan 5

zoneset name IVR_ZoneSet1
  zone name sample_vsan2-3
    pwwn 21:00:00:e0:8b:02:ca:4a vsan 3
    pwwn 21:00:00:20:37:c8:5c:6b vsan 2
```

Example 17-9 Displays Brief Information for an IVR VSAN Topology

```
switch# show ivr zoneset brief Active
zoneset name IVR_ZoneSet1
  zone name sample_vsan2-3
```

Example 17-10 Displays Brief Information for the Active IVZS

```
switch# show ivr zoneset brief Active
zoneset name IVR_ZoneSet1
  zone name sample_vsan2-3
```

Example 17-11 Displays Status Information for the IVZ

```
switch# show ivr zoneset status
Zoneset Status
```

| | |
|--------------------|----------------------------|
| name | : IVR_ZoneSet1 |
| state | : activation success |
| last activate time | : Sat Mar 22 21:38:46 1980 |
| force option | : off |

status per vsan:

| vsan | status |
|------|--------|
| 1 | active |
| 2 | active |

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 17-12 Displays the Specified Zone Set

```
switch# show ivr zoneset name IVR_ZoneSet1
zoneset name IVR_ZoneSet1
  zone name sample_vsan2-3
    pwwn 21:00:00:e0:8b:02:ca:4a vsan 3
    pwwn 21:00:00:20:37:c8:5c:6b vsan 2
```

Sample Configuration

This section provides the configuration steps to configure the example illustrated in [Figure 17-1](#).

Step 1 Enable IVR.

```
mds# config t
Enter configuration commands, one per line. End with CNTL/Z.
mds (config)# ivr enable
mds (config)# exit
```

Step 2 Verify that IVR is enabled.

```
mds# show ivr
Inter-VSAN Routing is enabled

Inter-VSAN enabled switches
-----
No IVR-enabled VSAN is active. Check VSAN-Topology configuration.

Inter-VSAN topology status
-----
Current Status: Inter-VSAN topology is INACTIVE

Inter-VSAN zoneset status
-----
      name           :
      state           : idle
      last activate time :
```

Step 3 Enable CFS distribution.

```
mds# config t
Enter configuration commands, one per line. End with CNTL/Z.
mds (config)# ivr distribution
mds (config)# exit
```

Send documentation comments to mdsfeedback-doc@cisco.com.

- Step 4** Manually configure the IVR VSAN-topology. In [Figure 17-1](#), two of the four IVR-enabled switches are members of VSANs 1 and 4. The other two switches are members of VSANs 2, 3, and 4.

```
mds# conf t
Enter configuration commands, one per line. End with CNTL/Z.
mds(config)# ivr vsan-topology database
mds(config-ivr-topology-db)# autonomous-fabric-id 1 switch-wwn 20:00:00:05:40:01:1b:c2
vsan-ranges 1,4
mds(config-ivr-topology-db)# autonomous-fabric-id 1 switch-wwn 20:02:00:44:22:00:4a:08
vsan-ranges 1,4
mds(config-ivr-topology-db)# autonomous-fabric-id 1 switch-wwn 20:00:00:44:22:02:8a:04
vsan-ranges 2-4
mds(config-ivr-topology-db)# autonomous-fabric-id 1 switch-wwn 20:00:00:44:22:40:aa:16
vsan-ranges 2-4
mds(config-ivr-topology-db)# exit
```

- Step 5** Verify the configured VSAN-topology.



Note The configured topology has not yet been activated—as indicated by the `no` status displayed in the Active column.

```
mds(config)# do show ivr vsan-topology
```

| AFID | SWITCH WWN | Active | Cfg. VSANS |
|------|---------------------------|--------|------------|
| 1 | 20:00:00:05:40:01:1b:c2 * | no | yes 1,4 |
| 1 | 20:00:00:44:22:00:4a:08 | no | yes 1,4 |
| 1 | 20:00:00:44:22:02:8a:04 | no | yes 2-4 |
| 1 | 20:00:00:44:22:40:aa:16 | no | yes 2-4 |

Total: 4 entries in active and configured IVR VSAN-Topology

Current Status: Inter-VSAN topology is INACTIVE

- Step 6** Activate the configured VSAN topology.

```
mds(config)# ivr vsan-topology activate
```

- Step 7** Verify the activation.

```
mds(config)# do show ivr vsan-topology
```

| AFID | SWITCH WWN | Active | Cfg. VSANS |
|------|---------------------------|--------|------------|
| 1 | 20:00:00:05:40:01:1b:c2 * | yes | yes 1,4 |
| 1 | 20:00:00:44:22:00:4a:08 | yes | yes 1,4 |
| 1 | 20:00:00:44:22:02:8a:04 | yes | yes 2-4 |
| 1 | 20:00:00:44:22:40:aa:16 | yes | yes 2-4 |

Total: 4 entries in active and configured IVR VSAN-Topology

Current Status: Inter-VSAN topology is ACTIVE

Last activation time: Tue May 20 23:14:59 1980

- Step 8** Configure IVR zone set and zones. Two zones are required:

- One zone has tape T (pwwn 10:02:50:45:32:20:7a:52) and server S1 (pwwn 10:02:66:45:00:20:89:04).
- Another zone has tape T and server S2 (pwwn 10:00:ad:51:78:33:f9:86).

Send documentation comments to mdsfeedback-doc@cisco.com.

**Tip**

Instead of creating two IVR zones, you can also create one IVR zone with the tape and both servers.

```
mds(config)# ivr zoneset name tape_server1_server2

mds(config-ivr-zoneset)# zone name tape_server1
mds(config-ivr-zoneset-zone)# member pwwn 10:02:50:45:32:20:7a:52 vsan 1
mds(config-ivr-zoneset-zone)# member pwwn 10:02:66:45:00:20:89:04 vsan 2
mds(config-ivr-zoneset-zone)# exit

mds(config-ivr-zoneset)# zone name tape_server2
mds(config-ivr-zoneset-zone)# member pwwn 10:02:50:45:32:20:7a:52 vsan 1
mds(config-ivr-zoneset-zone)# member pwwn 10:00:ad:51:78:33:f9:86 vsan 3
mds(config-ivr-zoneset-zone)# exit
```

- Step 9** View the IVR zone configuration to confirm that the IVR zone set and IVR zones are properly configured.

```
mds(config)# do show ivr zoneset
zoneset name tape_server1_server2
  zone name tape_server1
    pwwn 10:02:50:45:32:20:7a:52 vsan 1
    pwwn 10:02:66:45:00:20:89:04 vsan 2

  zone name tape_server2
    pwwn 10:02:50:45:32:20:7a:52 vsan 1
    pwwn 10:00:ad:51:78:33:f9:86 vsan 3
```

- Step 10** View the zone set prior to IVR zone set activation. Prior to activating the IVR zone set, view the active zone set. Repeat this step for VSANs 2 and 3.

```
mds(config)# do show zoneset active vsan 1
zoneset name finance_dept vsan 1
  zone name accounts_database vsan 1
    pwwn 10:00:23:11:ed:f6:23:12
    pwwn 10:00:56:43:11:56:fe:ee

  zone name $default_zone$ vsan 1
```

- Step 11** Activate the configured IVR zone set.

```
mds(config)# ivr zoneset activate name tape_server1_server2
zoneset activation initiated. check inter-VSAN zoneset status
mds(config)# exit
```

- Step 12** Verify the IVR zone set activation.

```
mds# show ivr zoneset active
zoneset name tape_server1_server2
  zone name tape_server1
    pwwn 10:02:50:45:32:20:7a:52 vsan 1
    pwwn 10:02:66:45:00:20:89:04 vsan 2

  zone name tape_server2
    pwwn 10:02:50:45:32:20:7a:52 vsan 1
    pwwn 10:00:ad:51:78:33:f9:86 vsan 3
```

Send documentation comments to mdsfeedback-doc@cisco.com.

- Step 13** Verify the zone set updates. Upon successful IVR zone set activation, verify that appropriate zones are added to the active zone set. Repeat this step for VSANs 2 and 3.

```

mds# show zoneset active vsan 1
zoneset name finance_dept vsan 1
  zone name accounts_database vsan 1
    pwwn 10:00:23:11:ed:f6:23:12
    pwwn 10:00:56:43:11:56:fe:ee

  zone name IVRZ_tape_server1 vsan 1
    pwwn 10:02:66:45:00:20:89:04
    pwwn 10:02:50:45:32:20:7a:52

  zone name IVRZ_tape_server2 vsan 1
    pwwn 10:02:50:45:32:20:7a:52
    pwwn 10:00:ad:51:78:33:f9:86

  zone name $default_zone$ vsan 1

mds# show ivr zoneset status
Zoneset Status
-----
name           : tape_server1_server2
state          : activation success
last activate time : Tue May 20 23:23:01 1980
force option    : on

status per vsan:
-----
vsan    status
-----
1       active

```

Default Settings

Table 17-5 lists the default settings for IVR parameters.

Table 17-5 **Default IVR Parameters**

| Parameters | Default |
|----------------------------|-------------------------------|
| IVR feature | Disabled. |
| IVR VSANs | Not added to virtual domains. |
| IVR NAT | Disabled. |
| QoS for IVZs | Low |
| Configuration Distribution | Disabled. |

Send documentation comments to mdsfeedback-doc@cisco.com.



Managing FLOGI, Name Server, FDMI, and RSCN Databases

This chapter describes the fabric login database, the name server features, the Fabric-Device Management Interface, and Registered State Change Notification (RSCN) information provided in the Cisco MDS 9000 Family. It includes the following sections:

- [Displaying FLOGI Details, page 18-1](#)
- [About the Name Server Proxy Feature, page 18-2](#)
- [Displaying FDMI, page 18-5](#)
- [About RSCN Information, page 18-7](#)

Displaying FLOGI Details

In a Fibre Channel fabric, each host or disk requires an FC ID. Use the **show flogi** command to verify if a storage device is displayed in the fabric login (FLOGI) table as in the following examples. If the required device is displayed in the FLOGI table, the fabric login is successful. Examine the FLOGI database on a switch that is directly connected to the host HBA and connected ports. See Examples [18-1](#) to [18-4](#).

Example 18-1 Displays Details on the FLOGI Database

```
switch# show flogi database
```

| INTERFACE | VSAN | FCID | PORT NAME | NODE NAME |
|-----------|------|----------|-------------------------|-------------------------|
| sup-fc0 | 2 | 0xb30100 | 10:00:00:05:30:00:49:63 | 20:00:00:05:30:00:49:5e |
| fc9/13 | 1 | 0xb200e2 | 21:00:00:04:cf:27:25:2c | 20:00:00:04:cf:27:25:2c |
| fc9/13 | 1 | 0xb200e1 | 21:00:00:04:cf:4c:18:61 | 20:00:00:04:cf:4c:18:61 |
| fc9/13 | 1 | 0xb200d1 | 21:00:00:04:cf:4c:18:64 | 20:00:00:04:cf:4c:18:64 |
| fc9/13 | 1 | 0xb200ce | 21:00:00:04:cf:4c:16:fb | 20:00:00:04:cf:4c:16:fb |
| fc9/13 | 1 | 0xb200cd | 21:00:00:04:cf:4c:18:f7 | 20:00:00:04:cf:4c:18:f7 |

Total number of flogi = 6.

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 18-2 Displays the FLOGI Database by Interface

```
switch# show flogi database interface fc1/11
```

| INTERFACE | VSAN | FCID | PORT NAME | NODE NAME |
|-----------|------|----------|-------------------------|-------------------------|
| fc1/11 | 1 | 0xa002ef | 21:00:00:20:37:18:17:d2 | 20:00:00:20:37:18:17:d2 |
| fc1/11 | 1 | 0xa002e8 | 21:00:00:20:37:38:a7:c1 | 20:00:00:20:37:38:a7:c1 |
| fc1/11 | 1 | 0xa002e4 | 21:00:00:20:37:6b:d7:18 | 20:00:00:20:37:6b:d7:18 |
| fc1/11 | 1 | 0xa002e2 | 21:00:00:20:37:18:d2:45 | 20:00:00:20:37:18:d2:45 |
| fc1/11 | 1 | 0xa002e1 | 21:00:00:20:37:39:90:6a | 20:00:00:20:37:39:90:6a |
| fc1/11 | 1 | 0xa002e0 | 21:00:00:20:37:36:0b:4d | 20:00:00:20:37:36:0b:4d |
| fc1/11 | 1 | 0xa002dc | 21:00:00:20:37:5a:5b:27 | 20:00:00:20:37:5a:5b:27 |
| fc1/11 | 1 | 0xa002da | 21:00:00:20:37:18:6f:90 | 20:00:00:20:37:18:6f:90 |
| fc1/11 | 1 | 0xa002d9 | 21:00:00:20:37:5b:cf:b9 | 20:00:00:20:37:5b:cf:b9 |
| fc1/11 | 1 | 0xa002d6 | 21:00:00:20:37:46:78:97 | 20:00:00:20:37:46:78:97 |

Total number of flogi = 10.

Example 18-3 Displays the FLOGI Database by VSAN

```
switch# show flogi database vsan 1
```

| INTERFACE | VSAN | FCID | PORT NAME | NODE NAME |
|-----------|------|----------|-------------------------|-------------------------|
| fc1/3 | 1 | 0xef02ef | 22:00:00:20:37:18:17:d2 | 20:00:00:20:37:18:17:d2 |
| fc1/3 | 1 | 0xef02e8 | 22:00:00:20:37:38:a7:c1 | 20:00:00:20:37:38:a7:c1 |
| fc1/3 | 1 | 0xef02e4 | 22:00:00:20:37:6b:d7:18 | 20:00:00:20:37:6b:d7:18 |
| fc1/3 | 1 | 0xef02e2 | 22:00:00:20:37:18:d2:45 | 20:00:00:20:37:18:d2:45 |
| fc1/3 | 1 | 0xef02e1 | 22:00:00:20:37:39:90:6a | 20:00:00:20:37:39:90:6a |
| fc1/3 | 1 | 0xef02e0 | 22:00:00:20:37:36:0b:4d | 20:00:00:20:37:36:0b:4d |
| fc1/3 | 1 | 0xef02dc | 22:00:00:20:37:5a:5b:27 | 20:00:00:20:37:5a:5b:27 |
| fc1/3 | 1 | 0xef02da | 22:00:00:20:37:18:6f:90 | 20:00:00:20:37:18:6f:90 |
| fc1/3 | 1 | 0xef02d9 | 22:00:00:20:37:5b:cf:b9 | 20:00:00:20:37:5b:cf:b9 |
| fc1/3 | 1 | 0xef02d6 | 22:00:00:20:37:46:78:97 | 20:00:00:20:37:46:78:97 |

Total number of flogi = 10.

Example 18-4 Displays the FLOGI Database by FC ID

```
switch# show flogi database fcid 0xef02e2
```

| INTERFACE | VSAN | FCID | PORT NAME | NODE NAME |
|-----------|------|----------|-------------------------|-------------------------|
| fc1/3 | 1 | 0xef02e2 | 22:00:00:20:37:18:d2:45 | 20:00:00:20:37:18:d2:45 |

Total number of flogi = 1.

See the “Default Company ID list” section on page 39-20 and the “Loop Monitoring Initiation” section on page 39-22.

About the Name Server Proxy Feature

The name server functionality maintains a database containing the attributes for all hosts and storage devices in each VSAN. Name servers allow a database entry to be modified by a device that originally registered the information.

The proxy feature is useful when you wish to modify (update or delete) the contents of a database entry that was previously registered by a different device.

Send documentation comments to mdsfeedback-doc@cisco.com.

All name server registration requests come from the same port whose parameter is registered or changed. If it does not, then the request is rejected.

This authorization enables WWNs to register specific parameters for another node.

Registering Name Server Proxies

To register the name server proxy, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fcns proxy-port 21:00:00:e0:8b:00:26:d0 vsan 2 | Configures a proxy port for the specified VSAN. |

Rejecting Duplicate pWWN

To prevent malicious or accidental log in using another device's pWWN, enable the **reject-duplicate-pwwn** option. If you disable this option, such pWWNs are allowed to log in to the fabric and replace the first device in the name server database.

To reject duplicate pWWNs, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fcns reject-duplicate-pwwn vsan 1 | Logs out devices when they log into the fabric if the pWWNs already exist. |
| | switch(config)# no fcns reject-duplicate-pwwn vsan 1 | Overwrites the first device's entry in the name server database with the new device having the same pwwn (default). |

Displaying Name Server Database Entries

The name server stores name entries for all hosts in the FCNS database. The name server permits an Nx port to register attributes during a PLOGI (to the name server) to obtain attributes of other hosts. These attributes are deregistered when the Nx port logs out either explicitly or implicitly.

In a multiswitch fabric configuration, the name server instances running on each switch shares information in a distributed database. One instance of the name server process runs on each switch.

Use the **show fcns** command to display the name server database and statistical information for a specified VSAN or for all VSANs (see Examples 18-5 to 18-8).

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 18-5 Displays the Name Server Database

```
switch# show fcns database
-----
FCID          TYPE  PWWN                               (VENDOR)          FC4-TYPE:FEATURE
-----
0x010000      N     50:06:0b:00:00:10:a7:80             (Cisco)            scsi-fcp fc-gs
0x010001      N     10:00:00:05:30:00:24:63             (Cisco)            ipfc
0x010002      N     50:06:04:82:c3:a0:98:52             (Company 1)        scsi-fcp 250
0x010100      N     21:00:00:e0:8b:02:99:36             (Company A)        scsi-fcp
0x020000      N     21:00:00:e0:8b:08:4b:20             (Company A)
0x020100      N     10:00:00:05:30:00:24:23             (Cisco)            ipfc
0x020200      N     21:01:00:e0:8b:22:99:36             (Company A)        scsi-fcp
```

Example 18-6 Displays the Name Server Database for the Specified VSAN

```
switch# show fcns database vsan 1
VSAN 1:
-----
FCID          TYPE  PWWN                               (VENDOR)          FC4-TYPE:FEATURE
-----
0x030001      N     10:00:00:05:30:00:25:a3             (Cisco)            ipfc
0x030101      NL    10:00:00:00:77:99:60:2c             (Interphase)
0x030200      N     10:00:00:49:c9:28:c7:01
0xec0001      NL    21:00:00:20:37:a6:be:14             (Seagate)          scsi-fcp
```

Total number of entries = 4

Example 18-7 Displays the Name Server Database Details

```
switch# show fcns database detail
-----
VSAN:1        FCID:0x030001
-----
port-wwn (vendor)      :10:00:00:05:30:00:25:a3 (Cisco)
node-wwn                :20:00:00:05:30:00:25:9e
class                   :2,3
node-ip-addr            :0.0.0.0
ipa                     :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:ipfc
symbolic-port-name      :
symbolic-node-name      :
port-type               :N
port-ip-addr            :0.0.0.0
fabric-port-wwn         :00:00:00:00:00:00:00:00
hard-addr               :0x000000
-----
VSAN:1        FCID:0xec0200
-----
port-wwn (vendor)      :10:00:00:5a:c9:28:c7:01
node-wwn                :10:00:00:5a:c9:28:c7:01
class                   :3
node-ip-addr            :0.0.0.0
ipa                     :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:
symbolic-port-name      :
symbolic-node-name      :
port-type               :N
port-ip-addr            :0.0.0.0
fabric-port-wwn         :22:0a:00:05:30:00:26:1e
hard-addr               :0x000000
Total number of entries = 2
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 18-8 Displays the Name Server Statistics

```
switch# show fcns statistics
registration requests received = 27
deregistration requests received = 0
queries received = 57
queries sent = 10
reject responses sent = 14
RSCNs received = 0
RSCNs sent = 0
```

Displaying FDMI

Cisco MDS SAN-OS Release 1.3 provides support for the Fabric-Device Management Interface (FDMI) functionally, as described in the FC-GS-4 standard. FDMI enables management of devices such as Fibre Channel Host Bus Adapters (HBAs) through in-band communications. This addition complements the existing Fibre Channel name server and management server functions.

Using the FDMI functionality, the SAN-OS software can extract the following management information about attached HBAs and host operating systems without installing proprietary host agents:

- Manufacturer, model, and serial number
- Node name and node symbolic name
- Hardware, driver, and firmware versions
- Host operating system (OS) name and version number

All FDMI entries are stored in persistent storage and are retrieved when the FDMI process is started.

Use the **show fdmi** command to display the FDMI database information (see Examples 18-9 to 18-11).

Example 18-9 Displays All HBA Management Servers

```
switch# show fdmi database
Registered HBA List for VSAN 1
  10:00:00:00:c9:32:8d:77
  21:01:00:e0:8b:2a:f6:54
switch# show fdmi database detail
Registered HBA List for VSAN 1
-----
HBA-ID: 10:00:00:00:c9:32:8d:77
-----
Node Name           :20:00:00:00:c9:32:8d:77
Manufacturer        :Emulex Corporation
Serial Num          :0000c9328d77
Model               :LP9002
Model Description:Emulex LightPulse LP9002 2 Gigabit PCI Fibre Channel Adapter
Hardware Ver        :2002606D
Driver Ver          :SLI-2 SW_DATE:Feb 27 2003, v5-2.20a12
ROM Ver             :3.11A0
Firmware Ver        :3.90A7
OS Name/Ver         :Window 2000
CT Payload Len      :1300000
Port-id: 10:00:00:00:c9:32:8d:77
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
-----
HBA-ID: 21:01:00:e0:8b:2a:f6:54
-----
Node Name           :20:01:00:e0:8b:2a:f6:54
Manufacturer        :QLogic Corporation
Serial Num          :\74262
Model               :QLA2342
Model Description:QLogic QLA2342 PCI Fibre Channel Adapter
Hardware Ver        :FC5010409-10
Driver Ver          :8.2.3.10 Beta 2 Test 1 DBG (W2K VI)
ROM Ver             :1.24
Firmware Ver        :03.02.13.
OS Name/Ver         :500
CT Payload Len      :2040
Port-id: 21:01:00:e0:8b:2a:f6:54
```

Example 18-10 Displays HBA Details for a Specified VSAN

```
switch# show fcsi database detail vsan 1
Registered HBA List for VSAN 1
-----
HBA-ID: 10:00:00:00:c9:32:8d:77
-----
Node Name           :20:00:00:00:c9:32:8d:77
Manufacturer        :Emulex Corporation
Serial Num          :0000c9328d77
Model               :LP9002
Model Description:Emulex LightPulse LP9002 2 Gigabit PCI Fibre Channel Adapter
Hardware Ver        :2002606D
Driver Ver          :SLI-2 SW_DATE:Feb 27 2003, v5-2.20a12
ROM Ver             :3.11A0
Firmware Ver        :3.90A7
OS Name/Ver         :Window 2000
CT Payload Len      :1300000
Port-id: 10:00:00:00:c9:32:8d:77
-----
HBA-ID: 21:01:00:e0:8b:2a:f6:54
-----
Node Name           :20:01:00:e0:8b:2a:f6:54
Manufacturer        :QLogic Corporation
Serial Num          :\74262
Model               :QLA2342
Model Description:QLogic QLA2342 PCI Fibre Channel Adapter
Hardware Ver        :FC5010409-10
Driver Ver          :8.2.3.10 Beta 2 Test 1 DBG (W2K VI)
ROM Ver             :1.24
Firmware Ver        :03.02.13.
OS Name/Ver         :500
CT Payload Len      :2040
Port-id: 21:01:00:e0:8b:2a:f6:54
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 18-11 Displays Details for the Specified HBA Entry

```
switch# show fdb database detail hba-id 21:01:00:e0:8b:2a:f6:54 vsan 1

Node Name           :20:01:00:e0:8b:2a:f6:54
Manufacturer        :QLogic Corporation
Serial Num          :\74262
Model               :QLA2342
Model Description:QLogic QLA2342 PCI Fibre Channel Adapter
Hardware Ver        :FC5010409-10
Driver Ver          :8.2.3.10 Beta 2 Test 1 DBG (W2K VI)
ROM Ver             :1.24
Firmware Ver        :03.02.13.
OS Name/Ver          :500
CT Payload Len      :2040
Port-id: 21:01:00:e0:8b:2a:f6:54
```

About RSCN Information

The Registered State Change Notification (RSCN) is a Fibre Channel service that informs hosts about changes in the fabric. Hosts can receive this information by registering with the fabric controller (through SCR). These notifications provide a timely indication of one or more of the following events:

- Disks joining or leaving the fabric.
- A name server registration change.
- A new zone enforcement.
- IP address change.
- Any other similar event that affects the operation of the host.

Apart from sending these events to registered hosts, a switch RSCN (SW-RSCN) is sent to all reachable switches in the fabric.



Note

The switch sends an RSCN to notify registered nodes that a change has occurred. It is up to the nodes to query the name server again to obtain the new information. The details of the changed information are not delivered by the switch in the RSCN sent to the nodes.

Displaying RSCN Information

Use the **show rscn** command to display RSCN information (see Examples 18-12 and 18-13).

Example 18-12 Displays Register Device Information

```
switch# show rscn scr-table vsan 1
SCR table for VSAN: 1
-----
FC-ID           REGISTERED FOR
-----
0x1b0300        fabric detected rscns
Total number of entries = 1
```

Send documentation comments to mdsfeedback-doc@cisco.com.



Note

The SCR table is not configurable. It is populated when hosts send SCR frames with RSCN information. If hosts do not receive RSCN information, then the **show rscn scr-table** command will not return entries.

Example 18-13 Displays RSCN Counter Information

```
switch# show rscn statistics vsan 1
Statistics for VSAN: 1
-----
Number of SCR received           = 8
Number of SCR ACC sent           = 8
Number of SCR RJT sent           = 0
Number of RSCN received          = 0
Number of RSCN sent              = 24
Number of RSCN ACC received      = 24
Number of RSCN ACC sent          = 0
Number of RSCN RJT received      = 0
Number of RSCN RJT sent          = 0
Number of SW-RSCN received       = 6
Number of SW-RSCN sent           = 15
Number of SW-RSCN ACC received   = 15
Number of SW-RSCN ACC sent       = 6
Number of SW-RSCN RJT received   = 0
Number of SW-RSCN RJT sent       = 0
```

About the multi-pid Option

If the RSCN **multi-pid** option is enabled, then RSCNs generated to the registered Nx ports may contain more than one affected port IDs. In this case, zoning rules are applied before putting the multiple affected port IDs together in a single RSCN. By enabling this option, you can reduce the number of RSCNs. For example: Suppose you have two disks (D1, D2) and a host (H) connected to switch 1. Host H is registered to receive RSCNs. D1, D2 and H belong to the same zone. If disks D1 and D2 are online at the same time, then one of the following applies:

- The **multi-pid** option is disabled on switch 1: two RSCNs are generated to host H—one for the disk D1 and another for disk D2.
- The **multi-pid** format is enabled on switch 1: a single RSCN is generated to host H, and the RSCN payload lists the affected port IDs (in this case, both D1 and D2).



Note

Some Nx ports may not understand multi-pid RSCN payloads. If so, disable the **multi-pid** RSCN option.

To configure the **multi-pid** option, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# rscn multi-pid vsan 105 | Sends RSCNs in a multi-pid format for VSAN 105. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Clearing RSCN Statistics

You can clear the counters and later view the counters for a different set of events. For example, you can keep track of how many RSCNs or SW-RSCNs are generated on a particular event (like ONLINE or OFFLINE events). You can use these statistics to monitor responses for each event in the VSAN.

Use the **clear rscn statistics** command to clear the RSCN statistics for the specified VSAN.

```
switch# clear rscn statistics vsan 1
```

After clearing the RSCN statistics, you can view the cleared counters by issuing the **show rscn** command.

```
switch# show rscn statistics vsan 1
Statistics for VSAN: 1
-----
Number of SCR received           = 0
Number of SCR ACC sent           = 0
Number of SCR RJT sent           = 0
Number of RSCN received          = 0
Number of RSCN sent              = 0
Number of RSCN ACC received      = 0
Number of RSCN ACC sent          = 0
Number of RSCN RJT received      = 0
Number of RSCN RJT sent          = 0
Number of SW-RSCN received       = 0
Number of SW-RSCN sent           = 0
Number of SW-RSCN ACC received   = 0
Number of SW-RSCN ACC sent       = 0
Number of SW-RSCN RJT received   = 0
Number of SW-RSCN RJT sent       = 0
```

Send documentation comments to mdsfeedback-doc@cisco.com.



Configuring Switch Security

The authentication, authorization, and accounting (AAA) mechanism verifies the identity of, grants access to, and tracks the actions of users managing a switch. All Cisco MDS 9000 Family switches use Remote Access Dial-In User Service (RADIUS) and Terminal Access Controller Access Control System Plus (TACACS+) protocols to provide solutions using remote AAA servers.

Based on the user ID and password combination provided, switches perform local authentication or authorization using the local database or remote authentication or authorization using AAA server(s). A preshared secret key provides security for communication between the switch and AAA servers. This secret key can be configured for all AAA server or for only a specific AAA server. This security mechanism provides a central management capability for AAA servers.

This chapter includes the following sections:

- [Switch Management Security, page 19-2](#)
- [Switch AAA Functionalities, page 19-2](#)
- [Configuring RADIUS, page 19-5](#)
- [Configuring TACACS+, page 19-10](#)
- [Configuring Server Groups, page 19-14](#)
- [Distributing AAA Server Configuration, page 19-15](#)
- [Local AAA Services, page 19-19](#)
- [Authentication and Authorization Process, page 19-20](#)
- [Role-Based Authorization, page 19-21](#)
- [Configuring User Accounts, page 19-29](#)
- [SNMP Security, page 19-32](#)
- [Configuring Accounting Services, page 19-32](#)
- [Configuring SSH Services, page 19-34](#)
- [Recovering Administrator Password, page 19-37](#)
- [Configuring Cisco ACS Servers, page 19-38](#)
- [Default Settings, page 19-42](#)

[Send documentation comments to mdsfeedback-doc@cisco.com.](mailto:mdsfeedback-doc@cisco.com)

Switch Management Security

Management security in any switch in the Cisco MDS 9000 Family provides security to all management access methods including the command-line interface (CLI) or Simple Network Management Protocol (SNMP).

CLI Security Options

You can access the CLI using the console (serial connection), Telnet, or Secure Shell (SSH). For each management path (console or Telnet and SSH), you can configure one or more of the following security control options: local, remote (RADIUS or TACACS+), or none.

- Remote security control
 - Using Remote Authentication Dial-In User Services (RADIUS). See the [“Configuring RADIUS” section on page 19-5](#).
 - Using Terminal Access Controller Access Control System plus (TACACS+). See the [“Configuring TACACS+” section on page 19-10](#).
- Local security control. See the [“Local AAA Services” section on page 19-19](#).

These security mechanisms can also be configured for the following scenarios:

- iSCSI authentication (see the [“Authentication Mechanism” section on page 28-71](#)).
- Fibre Channel Security Protocol (FC-SP) authentication (see the [Chapter 20, “Configuring Fabric Security”](#))

SNMP Security Options

The SNMP agent supports security features for SNMPv1, SNMPv 2c, and SNMPv3. Normal SNMP security mechanisms apply to all applications that use SNMP (for example, Cisco MDS 9000 Fabric Manager).

CLI security options also apply to the Cisco MDS Fabric Manager and Device Manager.

See [Chapter 22, “Configuring SNMP”](#).

Refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide* for information on the Cisco MDS Fabric or Device Managers.

Switch AAA Functionalities

Using the CLI or an SNMP application, you can configure authentication, authorization, and accounting (AAA) switch functionalities on any switch in the Cisco MDS 9000 Family.

Authentication

Authentication is the process of verifying the identity of the person managing the switch. This identity verification is based on the user ID and password combination provided by the person trying to manage the switch. Cisco MDS 9000 Family switches allow you to perform local authentication (using the local lookup database) or remote authentication (using one or more RADIUS or TACACS+ servers).

Send documentation comments to mdsfeedback-doc@cisco.com.

**Note**

When you log in to a Cisco MDS switch successfully using the Fabric Manager or Device Manager via Telnet or SSH and if that switch is configured for AAA server-based authentication, a temporary SNMP user entry is automatically created with an expiry time of one day. The SNMPv3 protocol data units (PDUs) with your Telnet/SSH login name as the SNMPv3 user are authenticated by the switch. The management station can temporarily use the Telnet/SSH login name as the SNMPv3 `auth` and `priv` passphrase. This temporary SNMP login is only allowed if you have one or more active MDS Shell sessions. If you do not have an active session at any given time, your login is deleted and you will not be allowed to perform SNMP v3 operations.

Authorization

By default, two roles exist in all Cisco MDS switches:

- Network operator (**`network-operator`**)—Has permission to view the configuration only. The operator cannot make any configuration changes.
- Network administrator (**`network-admin`**)—Has permission to execute all commands and make configuration changes. The administrator can also create and customize up to 64 additional roles.

If you use a SAN Volume Controller (SVC) setup, two more default roles exist in all Cisco MDS switches:

- SVC administrator (**`svc-admin`**)—Has permission to view the entire configuration and make SVC-specific configuration changes within the `switch(svc)` prompt.
- SVC operator (**`svc-operator`**)—Has permission to view the entire configuration. The operator cannot make any configuration changes.

**Note**

Refer to the *Cisco MDS 9000 Family SAN Volume Controller Configuration Guide* for more information on SVC.

These four default roles cannot be changed or deleted. You can create additional roles and configure the following options:

- Configure role-based authorization by assigning user roles locally or using remote AAA servers.
- Configure user profiles on a remote AAA server to contain role information. This role information is automatically downloaded and used when the user is authenticated through the remote AAA server.

**Note**

If a user only belongs to one of the newly-created roles and that role is subsequently deleted, then the user immediately defaults to the network-operator role.

Accounting

The accounting feature tracks and maintains a log of every management session used to access the switch. This information can be used to generate reports for troubleshooting and auditing purposes. Accounting logs can be stored locally or sent to remote AAA servers.

Send documentation comments to mdsfeedback-doc@cisco.com.

Remote AAA Services

Remote AAA services provided through RADIUS and TACACS+ protocols have the following advantages over AAA servers:

- It is easier to manage user password lists for each switch in the fabric.
- AAA servers are already deployed widely across enterprises and can be easily adopted.
- Easier to manage.
- Accounting log for all switches in the fabric can be centrally managed.
- Easier to manage user role mapping for each switch in the fabric.

Remote Authentication Guidelines

When you prefer using remote AAA servers, follow these guidelines:

- A minimum of one AAA server should be IP reachable.
- Be sure to configure a desired local AAA policy as this policy is used if all AAA servers are not reachable.
- AAA servers are easily reachable if an overlay Ethernet LAN is attached to the switch (see [Chapter 28, “Configuring IP Storage”](#)). This is the recommended method.
- SAN networks connected to the switch should have at least one gateway switch connected to the Ethernet LAN reaching the AAA servers.

Server Groups

You can specify remote AAA servers for authentication, authorization, and accounting using server groups. A server group is a set of remote AAA servers implementing the same AAA protocol. The purpose of a server group is to provide for fail-over servers in case a remote AAA server fails to respond. If the first remote server in the group fails to respond, the next remote server in the group is tried until one of the servers sends a response. If all the AAA servers in the server group fails to respond, then that server group option is considered a failure. If required, you can specify multiple server groups. If the Cisco MDS switch encounters errors from the servers in the first group, it tries the servers in the next server group.

You can create a server group using the **aaa group server** command.

AAA Service Configuration Options

AAA configuration in Cisco MDS 9000 Family switches is service based. You can have separate AAA configurations for the following services

- Telnet or SSH login (Cisco MDS Fabric Manager and Device Manager login).
- Console login.
- iSCSI authentication (see the [“Authentication Mechanism” section on page 28-71](#)).
- FC-SP authentication (see [Chapter 20, “Configuring Fabric Security”](#)).
- Accounting.

Send documentation comments to mdsfeedback-doc@cisco.com.

In general, server group, local, and none are the three options that can be specified for any service in an AAA configuration. Each option is tried in the order specified. If all the methods fail, local is tried.



Caution

Cisco MDS SAN-OS does not support all numeric usernames, whether created with TACACS+ or RADIUS, or created locally. Local users with all numeric names cannot be created. If an all numeric username exists on an AAA server and is entered during login, the user is not logged in.



Note

Even if local is not specified as one of the options, it is tried when all other configured options fail.

Table 19-1 provides the related CLI command for each AAA service configuration option.

Table 19-1 AAA Service Configuration Commands

| AAA Service Configuration Option | Related Command. |
|---|---|
| Telnet or SSH login (Cisco MDS Fabric Manager and Device Manager login) | Use the aaa authentication login default command. |
| Console login | Use the aaa authentication login console command. |
| iSCSI authentication | Use the aaa authentication iscsi default command. |
| FC-SP authentication | Use the aaa authentication dhchap default command. |
| Accounting | Use the aaa accounting default command |

Error-Enabled Status

When you log in, the login is processed by rolling over to local user database if the remote AAA servers do not respond. In such cases, the following message is displayed on the user's terminal—if you have enabled the error-enabled feature:

```
Remote AAA servers unreachable; local authentication done.
```

To enable this message display, use the **aaa authentication login error-enable** command.

To disable this message display, use the **no aaa authentication login error-enable** command.

To view the current display status, use the **show aaa authentication login error-enable** command (see [Example 19-6](#)).

Example 19-1 Displays AAA Authentication Login Information

```
switch# show aaa authentication login error-enable
enabled
```

Configuring RADIUS

Cisco MDS 9000 Family switches can use the RADIUS protocol to communicate with remote AAA servers. You can configure multiple RADIUS servers and server groups and set timeout and retry counts.

This section defines the RADIUS operation, identifies its network environments, and describes its configuration possibilities.

Send documentation comments to mdsfeedback-doc@cisco.com.

RADIUS is a distributed client/server protocol that secures networks against unauthorized access. In the Cisco implementation, RADIUS clients run on Cisco MDS 9000 Family switches and send authentication requests to a central RADIUS server that contains all user authentication and network service access information.

Setting the RADIUS Server Address

You can add up to 64 RADIUS servers. RADIUS keys are always stored in encrypted form in persistent storage. The running configuration also displays encrypted keys.

To specify the host RADIUS server address and the options, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# radius-server host 10.10.0.0 key HostKey | Specifies the preshared key for the selected RADIUS server. This key overrides the key assigned using the radius-server key command. In this example, the host is 10.10.0.0 and the key is HostKey. |
| Step 3 | switch(config)# radius-server host 10.10.0.0 auth-port 2003 | Specifies the destination UDP port number to which the RADIUS authentication messages should be sent. In this example, the host is 10.10.0.0 and the authentication port is 2003. The default authentication port is 1812, and the valid range is 0 to 65366. |
| Step 4 | switch(config)# radius-server host 10.10.0.0 acct-port 2004 | Specifies the destination UDP port number to which RADIUS accounting messages should be sent. The default accounting port is 1813, and the valid range is 0 to 65366. |
| Step 5 | switch(config)# radius-server host 10.10.0.0 accounting | Specifies this server to be used only for accounting purposes. Note If neither the authentication nor the accounting options are specified, the server is used for both accounting and authentication purposes. |
| Step 6 | switch(config)# radius-server host radius2 key 0 abcd | Specifies a clear text key for the specified server. The key is restricted to 64 characters. |
| | switch(config)# radius-server host radius3 key 7 da3Asda2ioyuoiuH | Specifies an encrypted key for the specified server. The key is restricted to 64 characters. |

Setting the Global Preshared Key

You need to configure the RADIUS preshared key to authenticate the switch to the RADIUS server. The length of the key is restricted to 64 characters and can include any printable ASCII characters (white spaces are not allowed). You can configure a global key to be used for all RADIUS server configurations on the switch.

You can override this global key assignment by explicitly using the **key** option in the **radius-server host** command.

Send documentation comments to mdsfeedback-doc@cisco.com.

To set the RADIUS preshared key, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# radius-server key AnyWord | Configures a preshared key (AnyWord) to authenticate communication between the RADIUS client and server. The default is clear text. |
| | switch(config)# radius-server key 0 AnyWord | Configures a preshared key (AnyWord) specified in clear text (indicated by 0) to authenticate communication between the RADIUS client and server. |
| | switch(config)# radius-server key 7 abe4DFeeweo00o | Configures a preshared key (specified in encrypted text) specified in encrypted text (indicated by 7) to authenticate communication between the RADIUS client and server. |

Setting the RADIUS Server Timeout Interval

To specify the time between retransmissions to the RADIUS servers, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# radius-server timeout 30 | Specifies the time (in seconds) between retransmissions to the RADIUS server. The default timeout is one (1) second. The time ranged from 1 to 60 seconds. |
| | switch(config)# no radius-server timeout 30 | Reverts the transmission time to its default (1) second. |

Setting Iterations of the RADIUS Server

By default, a switch retries a RADIUS server only once. This number can be configured. The maximum is five retries per server.

You can revert the retry number to its default by issuing the **no radius-server retransmit** command.

To specify the number of times that RADIUS servers should try to authenticate a user, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# radius-server retransmit 3 | Configures the number of times (3) the switch tries to connect to a RADIUS server(s) before reverting to local authentication. |

Defining Vendor-Specific Attributes

The Internet Engineering Task Force (IETF) draft standard specifies a method for communicating vendor-specific attributes (VSAs) between the network access server and the RADIUS server. The IETF uses attribute 26. VSAs allow vendors to support their own extended attributes that are not suitable for

Send documentation comments to mdsfeedback-doc@cisco.com.

general use. The Cisco RADIUS implementation supports one vendor-specific option using the format recommended in the specification. The Cisco vendor ID is 9, and the supported option is vendor type 1, which is named `cisco-avpair`. The value is a string with the following format:

```
protocol : attribute sep value *
```

Where `protocol` is a Cisco attribute for a particular type of authorization, and `sep` is = for mandatory attributes, and * is for optional attributes.

When you use RADIUS servers to authenticate yourself to a Cisco MDS 9000 Family switch, the RADIUS protocol directs the RADIUS server to return user attributes, such as authorization information, along with authentication results. This authorization information is specified through VSAs.

VSA Format

The following VSA protocol options are supported by the Cisco SAN-OS software:

- `Shell` protocol—used in access-accept packets to provide user profile information.
- `Accounting` protocol—used in accounting-request packets. If a value contains any white spaces, it should be put within double quotation marks.

The following attributes are supported by the Cisco SAN-OS software:

- `roles`—This attribute lists all the roles to which the user belongs. The value field is a string storing the list of group names delimited by white space. For example, if you belong to roles `vsan-admin` and `storage-admin`, the value field would be “`vsan-admin storage-admin`.” This subattribute is sent in the VSA portion of the Access-Accept frames from the RADIUS server, and it can only be used with the shell protocol value. These are two examples using the roles attribute:

```
shell:roles="network-admin vsan-admin"
```

```
shell:roles*"network-admin vsan-admin"
```

When an VSA is specified as `shell:roles*"network-admin vsan-admin"`, this VSA is flagged as an optional attribute, and other Cisco devices ignore this attribute.

- `accountinginfo`—This attribute stores additional accounting information besides the attributes covered by a standard RADIUS accounting protocol. This attribute is only sent in the VSA portion of the Account-Request frames from the RADIUS client on the switch, and it can only be used with the accounting protocol-related PDUs.

Specifying SNMPv3 on AAA Servers

The vendor/custom attribute `cisco-av-pair` can be used to specify user's role mapping using the format:

```
shell:roles="roleA roleB ..."
```

As of Cisco SAN-OS Release 2.0(1b), the VSA format is enhanced to optionally specify your SNMPv3 authentication and privacy protocol attributes also as follows:

```
shell:roles="roleA roleB..." snmpv3:auth=SHA priv=AES-128
```

The SNMPv3 authentication protocol options are SHA and MD5. The privacy protocol options are AES-128 and DES. If these options are not specified in the `cisco-av-pair` attribute on the ACS server, MD5 and DES are used by default.

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying RADIUS Server Details

Use the **show radius-server** command to display configured RADIUS parameters (see [Example 19-2](#)).

Example 19-2 Displays Configured RADIUS Information

```
switch# show radius-server
Global RADIUS shared secret:*****
retransmission count:5
timeout value:10
following RADIUS servers are configured:
  myradius.cisco.users.com:
    available for authentication on port:1812
    available for accounting on port:1813
  172.22.91.37:
    available for authentication on port:1812
    available for accounting on port:1813
    RADIUS shared secret:*****
  10.10.0.0:
    available for authentication on port:1812
    available for accounting on port:1813
    RADIUS shared secret:*****
```

Example 19-3 Displays Configured RADIUS Server-Group Order

```
switch# show radius-server groups
total number of groups:4
following RADIUS server groups are configured:
  group radius:
    server: all configured radius servers
  group Group1:
    server: Server3 on auth-port 1812, acct-port 1813
    server: Server5 on auth-port 1812, acct-port 1813
  group Group5:
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Configuring TACACS+

A Cisco MDS switch uses the Terminal Access Controller Access Control System Plus (TACACS+) protocol to communicate with remote AAA servers. You can configure multiple TACACS+ servers and set timeout values.

About TACACS+

TACACS+ is a client/server protocol that uses TCP (TCP port 49) for transport requirements. All switches in the Cisco MDS 9000 Family provide centralized authentication using the TACACS+ protocol. The addition of TACACS+ support in Cisco SAN-OS 1.3 enables the following advantages over RADIUS authentication:

- Provides independent, modular AAA facilities. Authorization can be done without authentication.
- TCP transport protocol to send data between the AAA client and server, using reliable transfers with a connection-oriented protocol.
- Encrypts the entire protocol payload between the switch and the AAA server to ensure higher data confidentiality. The RADIUS protocol only encrypts passwords.

Enabling TACACS+

By default, the TACACS+ feature is disabled in all switches in the Cisco MDS 9000 Family. You must explicitly enable the TACACS+ feature to access the configuration and verification commands for fabric authentication. When you disable this feature, all related configurations are automatically discarded.

To enable TACACS+ for a Cisco MDS switch, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# tacacs+ enable | Enables the TACACS+ in this switch. |
| | switch(config)# no tacacs+ enable | Disables (default) the TACACS+ in this switch. |

Setting the TACACS+ Server Address

If a secret key is not configured for a configured server, a warning message is issued if a global key is not configured. If a server key is not configured, the global key (if configured) is used for that server (see the [“Setting the Global Secret Key”](#) section on page 19-11).

Use the **tacacs-server** command to configure the communication parameters for the required TACACS+ server.

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure the TACACS+ server option, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# tacacs-server host 171.71.58.91 warning: no key is configured for the host | Configures the TACACS+ server identified by the specified IP address. |
| | switch(config)# no tacacs-server host 10.10.1.0 | Deletes the specified TACACS+ server identified by the IP address. By default, no server is configured. |
| Step 3 | switch(config)# tacacs-server host 171.71.58.91 port 2 | Configures the TCP port for all TACACS+ requests. |
| | switch(config)# no tacacs-server host 171.71.58.91 port 2 | Reverts to the factory default of using Port 49 for server access. |
| Step 4 | switch(config)# tacacs-server host host1.cisco.com key MyKey | Configures the TACACS+ server identified by the specified domain name and assigns the secret key. |
| Step 5 | switch(config)# tacacs-server host host100.cisco.com timeout 25 | Configures the timeout period for the switch to wait for a response from the specified server before it declares a timeout failure. |

Setting the Global Secret Key

You can configure global values for the **key** for all TACACS+ servers.



Note

If secret keys are configured for individual servers, those keys override the globally configured key.

To set the secret key for TACACS+ servers, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# tacacs-server key 7 3sdaA3daKUngd | Assigns the global secret key (in encrypted format) to access the TACACS+ server. This example specifies 7 to indicate the encrypted format being used. If this global key and the individual server keys are not configured, clear text messages are sent to the TACACS+ server(s). |
| | switch(config)# no tacacs-server key oldPword | Deletes the configured secret key to access the TACACS+ server and reverts to the factory default of allowing access to all configured servers. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Setting the Timeout Value

You can configure global timeout values for all TACACS+ servers.



Note

If timeout values are configured for individual servers, those values override the globally configured values.

To set the global timeout value for TACACS+ servers, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# tacacs-server timeout 30 | Configures the global timeout period for the switch to wait for a response from all servers before it declares a timeout failure. |
| | switch(config)# no tacacs-server timeout 30 | Deletes the configured timeout period and reverts to the factory default of 5 seconds. |

Defining Custom Attributes for Roles

Cisco MDS 9000 Family switches use the TACACS+ custom attribute for service shells to configure roles to which a user belongs. TACACS+ attributes are specified in `name=value` format. The attribute name for this custom attribute is `cisco-av-pair`. The following example illustrates how to specify roles using this attribute:

```
cisco-av-pair=shell:roles="network-admin vsan-admin"
```

You can also configure optional custom attributes to avoid conflicts with non-MDS Cisco switches using the same AAA servers.

```
cisco-av-pair*shell:roles="network-admin vsan-admin"
```

Additional custom attribute `shell:roles` are also supported:

```
shell:roles="network-admin vsan-admin"
```

or

```
shell:roles*"network-admin vsan-admin"
```



Note

TACACS+ custom attributes can be defined on an Access Control Server (ACS) for various services (for example, shell). Cisco MDS 9000 Family switches require the TACACS+ custom attribute for the service shell to be used for defining roles.

Send documentation comments to mdsfeedback-doc@cisco.com.

Supported TACACS+ Servers

The Cisco SAN-OS software currently supports the following parameters for the listed TACACS+ servers:

- TACACS:

```
cisco-av-pair=shell:roles="network-admin"
```

- Cisco ACS TACACS

```
shell:roles="network-admin"
shell:roles*"network-admin"
cisco-av-pair*shell:roles="network-admin"
cisco-av-pair*shell:roles*"network-admin"
cisco-av-pair=shell:roles*"network-admin"
```

- Open TACACS

```
cisco-av-pair*shell:roles="network-admin"
cisco-av-pair=shell:roles*"network-admin"
```

Displaying TACACS+ Server Details

Use the **show tacacs+** commands to display configurations for the TACACS+ protocol configuration in all switches in the Cisco MDS 9000 Family (see Examples 19-4 to 19-8).

Example 19-4 Displays Configured TACACS+ Server Information

```
switch# show tacacs-server
Global TACACS+ shared secret:*****
timeout value:30
total number of servers:3

following TACACS+ servers are configured:
  171.71.58.91:
    available on port:2
  cisco.com:
    available on port:49
  171.71.22.95:
    available on port:49
    TACACS+ shared secret:*****
```

Example 19-5 Displays AAA Authentication Information

```
switch# show aaa authentication
default: group TacServer local none
console: local
iscsi: local
dhchap: local
```

Example 19-6 Displays AAA Authentication Login Information

```
switch# show aaa authentication login error-enable
enabled
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 19-7 Displays Configured TACACS Server Groups

```
switch# show tacacs-server groups
total number of groups:2

following TACACS+ server groups are configured:
  group TacServer:
    server 171.71.58.91 on port 2
  group TacacsServer1:
    server ServerA on port 49
    server ServerB on port 49:
```

Example 19-8 Displays All AAA Server Groups

```
switch# show aaa groups
radius
TacServer
```

Configuring Server Groups

You can specify one or more remote AAA servers to authenticate users using server groups. All members of a group must belong to the same protocol: either RADIUS or TACACS+. The servers are tried in the same order in which you configure them.

You can configure these server groups at any time but they only take effect when you apply them to a AAA service.

Use the **aaa authentication** command to apply server groups to a AAA service. For example, you can use the **aaa authentication** command to configure AAA policies for CLI or Fabric Manager or Device Manager users.

To specify the TACACS+ server order within a group, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# aaa group server tacacs+ TacacsServer1 switch(config-tacacs+)# | Configures a TacacsServer1 group and enters the submode for that group. |
| | switch(config)# no aaa group server tacacs+ TacacsServer19 | Deletes the group called TacacsServer19 from the authentication list. |
| Step 3 | switch(config-tacacs+)# server ServerA | Configures ServerA to be tried first within the server group called the TacacsServer1. Tip If the specified TACACS+ server is not found, configure it using the tacacs-server host command and retry this command. |
| Step 4 | switch(config-tacacs+)# server ServerB | Configures ServerB to be tried second within TacacsServer1. |
| | switch(config-tacacs+)# no server ServerZ | Deletes ServerZ within the TacacsServer1 list of servers. |

Send documentation comments to mdsfeedback-doc@cisco.com.

To verify the configured server group order, use the **show tacacs-server groups** command:

```
switch# show tacacs-server groups
total number of groups:2

following TACACS+ server groups are configured:
  group TacServer:
    server 171.71.58.91 on port 2
  group TacacsServer1:
    server ServerA on port 49
    server ServerB on port 49:
```

Distributing AAA Server Configuration

Configuration for RADIUS and TACACS+ AAA on a MDS switch can be distributed using the Cisco Fabric Services (CFS). The distribution is disabled by default (see [Chapter 9, “Using the CFS Infrastructure”](#)).

After enabling the distribution, the first server or global configuration starts an implicit session. All server configuration commands entered thereafter are stored in a temporary database and applied to all switches in the fabric (including the originating one) when you explicitly commit the database. The various server and global parameters are distributed, except the server and global keys. These keys are unique secrets to a switch and should not be shared with other switches.



Note

Server group configurations are not distributed.

Enabling the RADIUS Server Distribution

Only switches where distribution is enabled can participate in the distribution activity.

To enable RADIUS server distribution, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# radius distribute | Enables RADIUS configuration distribution in this switch. |
| | switch(config)# no radius distribute | Disables RADIUS configuration distribution in this switch (default). |

To enable TACACS+ server distribution, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# tacacs+ distribute | Enables TACACS+ configuration distribution in this switch. |
| | switch(config)# no tacacs+ distribute | Disables TACACS+ configuration distribution in this switch. (default) |

Send documentation comments to mdsfeedback-doc@cisco.com.

Starting a Distribution Session on a Switch

A distribution session starts the moment you begin a RADIUS/TACACS+ server or global configuration. For example, the following tasks start an implicit session:

- Specifying the global timeout for RADIUS servers.
- Specifying the global timeout for TACACS+ servers.



Note

After you issue the first configuration command related to AAA servers, all server and global configurations made (including the configuration that caused the distribution session start) are stored in a temporary buffer—not in the running configuration.

To specify the global timeout and start an implicit session for RADIUS servers, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# radius-server timeout 30 | Specifies the time (in seconds) between retransmissions to the RADIUS server. The default timeout is one (1) second. The time range in seconds is 1 to 60. |

To specify the global timeout and start an implicit session for TACACS+ servers, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# tacacs-server timeout 30 | Configures the global timeout period for the switch to wait for a response from all servers before it declares a timeout failure. |
| | switch(config)# no tacacs-server timeout 30 | Deletes the configured timeout period and reverts to the factory default of 5 seconds. |

Displaying the Session Status

Once the implicit distribution session has started, you can check the session status using the **show radius distribution status** command.

```
switch# show radius distribution status
distribution : enabled
session ongoing: yes
session owner: admin
session db: exists
merge protocol status: merge activation done
```

```
last operation: enable
last operation status: success
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Once the implicit distribution session has started, you can check the session status using the **show tacacs distribution status** command.

```
switch# show tacacs+ distribution status
distribution : enabled
session ongoing: yes
session owner: admin
session db: exists
merge protocol status: merge activation done

last operation: enable
last operation status: success
```

Displaying the Configuration to Be Distributed

To display the RADIUS global and/or server configuration stored in the temporary buffer, use the **show radius pending** command.

```
switch(config)# show radius pending-diff
+radius-server host testhost1 authentication accounting
+radius-server host testhost2 authentication accounting
```

To display the TACACS global and/or server configuration stored in the temporary buffer, use the **show tacacs+ pending** command.

```
switch(config)# show tacacs+ pending-diff
+tacacs-server host testhost3
+tacacs-server host testhost4
```

Committing the Distribution

The RADIUS or TACACS global and/or server configuration stored in the temporary buffer can be applied to the running configuration across all switches in the fabric (including the originating switch).

To commit RADIUS distribution, use the **radius commit** command, to commit TACACS distribution, use the **tacacs+ commit** command:

Discarding the Distribution Session

Discarding the distribution of a session-in-progress causes the configuration in the temporary buffer to be dropped. The distribution is no applied.

To discard the RADIUS session-in-progress distribution, use the **radius abort** command, and to discard the TACACS+ session-in-progress distribution, use the **tacacs+ abort** command.

Clearing Sessions

To clear the ongoing CFS distribution session (if any) and to unlock the fabric for the RADIUS feature, issue the **clear radius session** command from any switch in the fabric.

```
switch# clear radius session
```

Send documentation comments to mdsfeedback-doc@cisco.com.

To clear the ongoing CFS distribution session (if any) and to unlock the fabric for the TACACS+ feature, issue the **clear tacacs+ session** command from any switch in the fabric.

```
switch# clear tacacs+ session
```

Merge Guidelines for RADIUS and TACACS+ Configurations

The RADIUS and TACACS+ server and global configuration are merged when two fabrics merge. The merged configuration is applied to CFS distribution-enabled switches.

When merging the fabric, be aware of the following conditions:

- The server groups are not merged
- The server and global keys are not changed during merge
- The merged configuration contains all servers found on all CFS enabled switches
- The timeout and retransmit parameters of the merged configuration are the largest values found per server and global.



Caution

If there is a conflict between two switches in the server ports configured, the merge fails.

Use the **show radius distribution status** command to view the status of the RADIUS fabric merge (see [Example 19-9](#)).

Example 19-9 Displays the RADIUS Fabric Merge Status

```
switch# show radius distribution status
distribution : enabled
session ongoing: no
session db: does not exist
merge protocol status: merge response received
merge error: conflict: server dmttest2 has auth-port 1812 on this switch and 1999
on remote

last operation: enable
last operation status: success
```

Use the **show tacacs+ distribution status** command to view the status of the TACACS+ fabric merge (see [Example 19-10](#)).

Example 19-10 Displays the TACACS+ Fabric Merge Status

```
switch# show tacacs+ distribution status
distribution : enabled
session ongoing: no
session db: does not exist
merge protocol status: merge activation done

last operation: enable
last operation status: success
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Local AAA Services

The system maintains the user name and password locally and stores the password information in encrypted form. You are authenticated based on the locally stored user information.

Use the **username** command to configure local users and their roles (see the “[Creating or Updating Users](#)” section on page 19-29).

Use the **show accounting log** command to view the local accounting log (see [Example 19-11](#)).

Example 19-11 Displays the Accounting Log Information

```
switch# show accounting log

Sat Jan 24 03:22:06 1981:stop:snmp_349154526_171.71.58.69:admin:
Sat Jan 24 03:22:06 1981:start:snmp_349154526_171.71.58.69:admin:
Sat Jan 24 03:22:06 1981:update:snmp_349154526_171.71.58.69:admin:Added member [
  WWN: 21:00:00:20:37:a6:be:00 ID: 2] to zone test-27 on VSAN 1
...
Sat Jan 24 23:59:56 1981:stop:/dev/pts/0_349228792:root:shell terminated
Sun Jan 25 00:00:06 1981:start:/dev/pts/1_349228806:admin:
```

Disabling AAA Authentication

You can turn off password verification using the **none** option. If you configure this option, users can login without giving a valid password. But the user should at least exist locally on the Cisco MDS 9000 Family switch.



Caution

Use this option cautiously. If configured, any user will be able to access the switch at any time.

Use the **none** option in the **aaa authentication login** command to disable password verification.

A user created using the **username** command will exist locally on the Cisco MDS 9000 Family switch.

Displaying AAA Authentication

The **show aaa authentication** command displays the configured authentication methods (see [Example 19-12](#)).

Example 19-12 Example 16-8 Displays Authentication Information

```
switch# show aaa authentication

No AAA Authentication
default: group TacServer local none
console: local none
iscsi: local
dhchap: local
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Authentication and Authorization Process

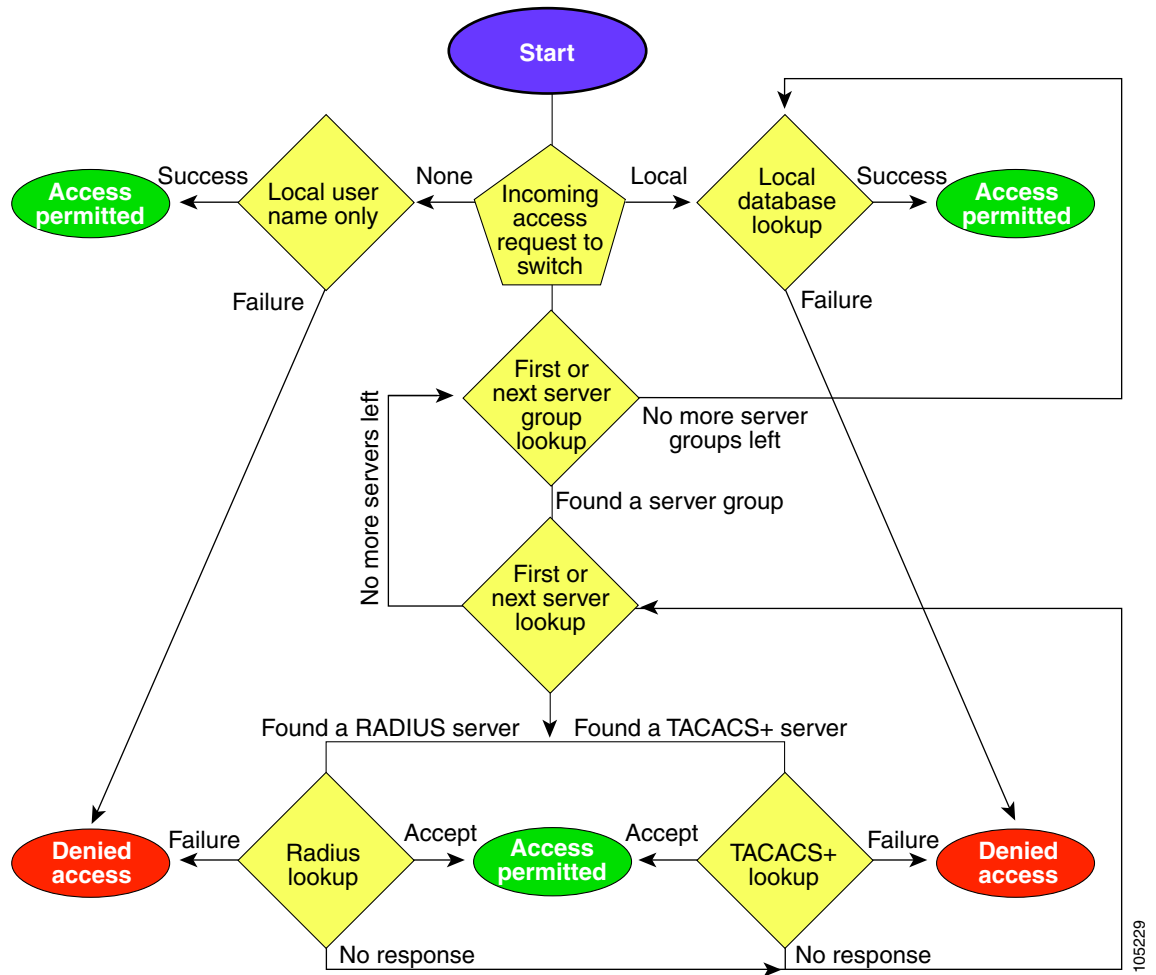
Authentication is the process of verifying the identity of the person managing the switch. This identity verification is based on the user ID and password combination provided by the person trying to manage the switch. The Cisco MDS 9000 Family switches allow you to perform local authentication (using the lookup database) or remote authentication (using one or more RADIUS servers or TACACS+ servers).

Figure 19-1 shows a flow chart of the process. The following steps explain the authorization and authentication process.

-
- Step 1** When you can log in to the required switch in the Cisco MDS 9000 Family, you can use the Telnet, SSH, Fabric Manager/Device Manager, or console login options.
- Step 2** When you have configured server groups using the server group authentication method, an authentication request is sent to the first AAA server in the group.
- If the AAA server fails to respond, then the next AAA server is tried and so on until the remote server responds to the authentication request.
 - If all AAA servers in the server group fail to respond, then the servers in the next server group are tried.
 - If all configured methods fail, then the local database is used for authentication.
- Step 3** If you are successfully authenticated through a remote AAA server, then the following possibilities apply.
- If AAA server protocol is RADIUS, then user roles specified in the `cisco-av-pair` attribute are downloaded with an authentication response.
 - If AAA server protocol is TACACS+, then another request is sent to the same server to get the user roles specified as custom attributes for the shell.
 - If user roles are not successfully retrieved from the remote AAA server, then the user is assigned the network-operator role.
- Step 4** If your user name and password are successfully authenticated locally, you are allowed to log in, and you are assigned the roles configured in the local database.
-

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 19-1 Switch Authorization and Authentication Flow



Note

No more server groups left = no response from any server in all server groups.
No more servers left = no response from any server within this server group.



Tip

In Step 1, use the **aaa authentication login default** command to configure policies for using Telnet, SSH, or Fabric Manager/Device Manager and the **aaa authentication login console** command to configure AAA policies using the console. If the **aaa authentication login console** command is not configured for console login, the software automatically uses policies used by the **aaa authentication login default** command.

Role-Based Authorization

Switches in the Cisco MDS 9000 Family perform authentication based on roles. Role-based authorization limits access to switch operations by assigning users to roles. This kind of authentication restricts you to management operations based on the roles to which you have been added.

Send documentation comments to mdsfeedback-doc@cisco.com.

When you execute a command, perform command completion, or obtain context sensitive help, the switch software allows the operation to progress if you have permission to access that command.

Each role can contain multiple users and each user can be part of multiple roles. For example, if role1 users are only allowed access to configuration commands, and role2 users are only allowed access to debug commands, then if Joe belongs to both role1 and role2, he can access configuration as well as debug commands.



Note

If you belong to multiple roles, you can execute a union of all the commands permitted by these roles. Access to a command takes priority over being denied access to a command. For example, suppose you belong to a TechDocs group and you were denied access to configuration commands. However, you also belong to the engineering group and have access to configuration commands. In this case, you will have access to configuration commands.



Tip

Any role, when created, does not allow access to the required commands immediately. The administrator must configure appropriate rules for each role to allow access to the required commands.

Configuring Roles and Profiles

To create an additional role or to modify the profile for an existing role, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# role name techdocs switch(config-role)# | Places you in the mode for the specified role (techdocs). Note The role submode prompt indicates that you are now in the role submode. This submode is now specific to the techdocs group. |
| | switch(config)# no role name techdocs | Deletes the role called techdocs. |
| Step 3 | switch(config-role)# description Entire Tech. Docs. group | Assigns a description to the new role. The description is limited to one line and can contain spaces. |
| | switch(config-role)# no description | Resets the description for the Tech. Docs. group. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Configuring Rules and Features for Each Role

Up to 16 rules can be configured for each role. The user-specified rule number determines the order in which the rules are applied. For example, rule 1 is applied before rule 2, which is applied before rule 3, and so on. A user not belonging to the network-admin role cannot perform commands related to roles.

For example, if user A is permitted to perform all **show** commands, user A cannot view the output of the **show role** command if user A does not belong to the network-admin role

The **rule** command specifies operations that can be performed by a specific role. Each rule consists of a rule number, a rule type (permit or deny), a command type (for example, **config**, **clear**, **show**, **exec**, **debug**), and an optional feature name (for example, FSPF, zone, VSAN, fcping, or interface).



Note

In this case, **exec** commands refer to all commands in the EXEC mode that do not fall in the **show**, **debug**, and **clear**, categories.

Modifying Profiles

To modify the profile for an existing role, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# role name sangroup switch(config-role)# | Places you in sangroup role submenu. |
| Step 3 | switch(config-role)# rule 1 permit config switch(config-role)# rule 2 deny config feature fspf switch(config-role)# rule 3 permit debug feature zone switch(config-role)# rule 4 permit exec feature fcping | Allows users belonging to the sangroup role to perform all configuration commands except fspf config commands. They can also perform zone debug commands and the fcping EXEC mode command. |
| Step 4 | switch(config-role)# no rule 4 | Deletes rule 4, which no longer permits the sangroup to perform the fcping command. |

In Step 3, rule 1 is applied first, thus permitting sangroup users access to all **config** commands. Rule 2 is applied next, denying FSPF configuration to sangroup users. As a result, sangroup users can perform all other **config** commands, except **fspf** configuration commands.



Note

The order of rule placement is important. If you had swapped these two rules and issued the **deny config feature fspf** rule first and issued the **permit config** rule next, you would be allowing all sangroup users to perform all configuration commands because the second rule globally overrode the first rule.

Send documentation comments to mdsfeedback-doc@cisco.com.

Configuring the VSAN Policy

Configuring the VSAN policy requires the ENTERPRISE_PKG license (see [Chapter 3, “Obtaining and Installing Licenses”](#)).

You can configure a role so that it only allows tasks to be performed for a selected set of VSANs. By default, the VSAN policy for any role is **permit**, which allows tasks to be performed for all VSANs. You can configure a role that only allows tasks to be performed for a selected set of VSANs. To selectively allow VSANs for a role, set the VSAN policy to **deny**, and then set the configuration to **permit** or the appropriate VSANs.



Note

Users configured in roles where the VSAN policy is set to **deny** cannot modify the configuration for E ports. They can only modify the configuration for F or FL ports (depending on whether the configured rules allow such configuration to be made). This is to prevent such users from modifying configurations that may impact the core topology of the fabric.



Tip

Roles can be used to create VSAN administrators. Depending on the configured rules, these VSAN administrators can configure MDS features (for example, zone, fcdomain, or VSAN properties) for their VSANs without affecting other VSANs. Also, if the role permits operations in multiple VSANs, then the VSAN administrators can change VSAN membership of F or FL ports among these VSANs.

Users belonging to roles in which the VSAN policy is set to **deny** are referred to as VSAN-restricted users. These users cannot perform commands that require the startup configuration to be viewed or modified.

These commands include the **copy running-config startup-config**, **show startup-config**, **show running-config diff**, and **copy startup-config running-config** commands. For information on these commands, see [Chapter 2, “Before You Begin.”](#)

Modifying the VSAN Policy

To modify the VSAN policy for an existing role, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# role name sangroup switch(config-role)# | Places you in sangroup role submode. |
| Step 3 | switch(config)# vsan policy deny switch(config-role-vsan) | Changes the VSAN policy of this role to deny and places you in a submode where VSANs can be selectively permitted. |
| Step 4 | switch(config-role)# no vsan policy deny | Deletes the configured VSAN role policy and reverts to the factory default (permit). |
| | switch(config-role-vsan)# permit vsan 10-30 | Permits this role to perform the allowed commands for VSANs 10 through 30. |
| | switch(config-role-vsan)# no permit vsan 15-20 | Removes the permission for this role to perform commands for vsan 15 to 20. So, the role is now permitted to perform commands for VSAN 10 to 14, and 21 to 30. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Distributing Role-Based Configurations

Role-based configurations use the Cisco Fabric Services (CFS) infrastructure to enable efficient database management, provide a single point of configuration for the entire fabric (see [Chapter 9, “Using the CFS Infrastructure”](#)).

The following configurations are distributed:

- Role names and descriptions
- List of rules for the roles
- VSAN policy and the list of permitted VSANs

Database Implementation

Role-based configurations uses two databases to accept and implement configurations.

- Configuration database—The database currently enforced by the fabric.
- Pending database—Your subsequent configuration changes are stored in the pending database. If you modify the configuration, you need to commit or discard the pending database changes to the configuration database. The fabric remains locked during this period. Changes to the pending database are not reflected in the configuration database until you commit the changes.

Locking The Fabric

The first action that modifies the database creates the pending database and locks the feature in the entire fabric. Once you lock the fabric, the following situations apply:

- No other user can make any configuration changes to this feature.
- A copy of the configuration database becomes the pending database along with the first change.

Committing the Changes

If you commit the changes made to the pending database, the configuration is committed to all the switches in the fabric. On a successful commit, the configuration change is applied throughout the fabric and the lock is released. The configuration database now contains the committed changes and the pending database is now cleared.

To commit role-based configuration changes, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# role commit vsan 3 | Commits the role-based configuration changes. |

Discarding the Changes

If you discard (abort) the changes made to the pending database, the configuration database remains unaffected and the lock is released.

Send documentation comments to mdsfeedback-doc@cisco.com.

To discard role-based configuration changes, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# role abort | Discards the role-based configuration changes and clears the pending configuration database. |

Enabling Distribution

To enable role-based configuration distribution, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# role distribute | Enables role-based configuration distribution. |
| | switch(config)# no role distribute | Disables role-based configuration distribution (default). |

Clearing Sessions

To forcibly clear the existing role session in the fabric, issue the **clear role session** command from any switch that is part of the initiated session.



Caution

Any changes in the pending database are lost when you issue this command.

```
switch# clear role session
```

Database Merge Guidelines

Fabric merge does not modify the role database on a switch. If two fabrics merge, and the fabrics have different role databases, the software generates an alert message.

Refer to the “[CFS Merge Support](#)” section on page 9-7 for detailed concepts.

- Verify that the role database is identical on all switches in the entire fabric
- Be sure to edit the role database on any switch to the desired database and then commit it. This synchronizes the role databases on all the switches in the fabric.

Displaying Role-Based Information

Use the **show role** command to display rules configured on the switch. The rules are displayed by rule number and are based on each role. All roles are displayed if the role name is not specified. See [Example 19-13](#).

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 19-13 Displays Information for All Roles

```
switch# show role
Role: network-admin
Description: Predefined Network Admin group. This role cannot be modified
Access to all the switch commands

Role: network-operator
Description: Predefined Network Operator group. This role cannot be modified
Access to Show commands and selected Exec commands

Role: svc-admin
Description: Predefined SVC Admin group. This role cannot be modified
Access to all SAN Volume Controller commands

Role: svc-operator
Description: Predefined SVC Operator group. This role cannot be modified
Access to selected SAN Volume Controller commands

Role: TechDocs
    vsan policy: permit (default)

Role: sangroup
    Description: SAN management group
    vsan policy: deny
    Permitted vsans: 10-30
```

| Rule | Type | Command-type | Feature |
|------|--------|--------------|---------|
| 1. | permit | config | * |
| 2. | deny | config | fspf |
| 3. | permit | debug | zone |
| 4. | permit | exec | fcping |

Displaying Role-Based When Distribution is Enabled

Use the **show role** command to display the configuration database.

Use the **show role status** command to display whether distribution is enabled for role configuration, the current fabric status (locked or unlocked) and the last operation performed. See [Example 19-14](#).

Example 19-14 Displays the Role Status Information

```
switch# show role status
Distribution: Enabled
Session State: Locked

Last operation (initiated from this switch): Distribution enable
Last operation status: Success
```

Use the **show role pending** command to display the pending role database. See [Example 19-15](#).

[Example 19-15](#) displays the output of the **show role pending** command if you follow the following procedure:

- Create the role called `myrole` using the **role name myrole** command.
- Issue the **rule 1 permit config feature fspf** command.
- Issue the **show role pending** command to see the output in [Example 19-15](#).

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 19-15 Displays Information on the Pending Roles Database

```
switch# show role pending
Role: network-admin
Description: Predefined Network Admin group. This role cannot be modified
Access to all the switch commands

Role: network-operator
Description: Predefined Network Operator group. This role cannot be modified
Access to Show commands and selected Exec commands

Role: svc-admin
Description: Predefined SVC Admin group. This role cannot be modified
Access to all SAN Volume Controller commands

Role: svc-operator
Description: Predefined SVC Operator group. This role cannot be modified
Access to selected SAN Volume Controller commands

Role: TechDocs
  vsan policy: permit (default)

Role: sangroup
  Description: SAN management group
  vsan policy: deny
  Permitted vsans: 10-30
```

| Rule | Type | Command-type | Feature |
|------|--------|--------------|---------|
| 1. | permit | config | * |
| 2. | deny | config | fspf |
| 3. | permit | debug | zone |
| 4. | permit | exec | fcping |

```
Role: myrole
  vsan policy: permit (default)

-----
Rule      Type      Command-type      Feature
-----
1.  permit      config              fspf
```

Use the **show role pending-diff** command to display the differences between the pending and configuration role database. See [Example 19-16](#).

Example 19-16 Displays the Differences between the Two Databases

```
switch# show role pending-diff
+Role: myrole
+  vsan policy: permit (default)
+  -----
+  Rule      Type      Command-type      Feature
+  -----
+  1.  permit      config              fspf
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Configuring User Accounts

Every Cisco MDS 9000 Family switch user has the account information stored by the system. Your authentication information, user name, user password, password expiration date, and role membership are stored in your user profile.

The tasks explained in this section enable you to create users and modify the profile of an existing user. These tasks are restricted to privileged users as determined by your administrator.

The password should have the strong characteristics, such as the following:

- Are at least eight characters long
- Not contain many consecutive characters (such as “abcd”)
- Not contain many repeating characters (such as “aaabbb”)
- Not contain dictionary words
- Contain both upper and lower case characters
- Contain numbers

Creating or Updating Users

As of Cisco SAN-OS Release 2.0(1b), the passphrase specified in the **snmp-server user** option and the password specified **username** option are synchronized (see the “[SNMPv3 CLI User Management and AAA Integration](#)” section on page 22-3).

By default, the user account does not expire unless you explicitly configure it to expire. The **expire** option determines the date on which the user account is disabled. The date is specified in the YYYY-MM-DD format.



Tip

The following words are reserved and cannot be used to configure users: bin, daemon, adm, lp, sync, shutdown, halt, mail, news, uucp, operator, games, gopher, ftp, nobody, nsd, mailnull, rpc, rpcuser, xfs, gdm, mtsuser, ftpuser, man, and sys.



Note

User passwords are not displayed in the switch configuration file.



Tip

If a password is trivial (short, easy-to-decipher), your password configuration is rejected. Be sure to configure a strong password as shown in the sample configuration. Passwords are case-sensitive. As of Release 2.0(1b), “admin” is no longer the default password for any Cisco MDS 9000 Family switch. You must explicitly configure a strong password.



Caution

Cisco MDS SAN-OS does not support all numeric usernames, whether created with TACACS+ or RADIUS, or created locally. Local users with all numeric names cannot be created. If an all numeric username exists on an AAA server and is entered during login, the user is not logged in.

Send documentation comments to mdsfeedback-doc@cisco.com.

**Tip**

To issue commands with the **internal** keyword for troubleshooting purposes, you must have an account that is a member of the network-admin group.

To configure a new user or to modify the profile of an existing user, follow these steps:

| | Command | Purpose |
|---------------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# username usam password abcd123AAA expire 2003-05-31 | Creates or updates the user account (usam) along with a password (abcd123AAA) that is set to expire on 2003-05-31. The password is limited to 64 characters. Note User account names must contain non-numeric characters. |
| | switch(config)# username msam password 0 abcd12AAA role network-operator | Creates or updates the user account (msam) along with a password (abcd12AAA) specified in clear text (indicated by 0). The password is limited to 64 characters. Note User account names must contain non-numeric characters. |
| | switch(config)# username user1 password 5 !*asdfsdfjh!@df | Specifies an encrypted (specified by 5) password (!@*asdfsdfjh!@df) for the user account (user1). |
| Step 3 | switch(config)# username usam role network-admin | Adds the specified user (usam) to the network-admin role. |
| | switch(config)# no username usam role vsan-admin | Deletes the specified user (usam) from the vsan-admin role. |
| Step 4 | switch(config)# username admin sshkey ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEAtjIHrIt/3dDeohix6JcRSIYZ 0EOdJ3l5RONWcwSgAuTUSrLk 3a9hdYkzY94fhHmNGQGCjVg+8cbOxyH4Z1jcVFcrDogtQT+Q8dve qts/8XQhqkNAFeGy4u8TJ2Us oreCU6DlibwkpzDafzKTPA5vB6FmHd2TI6Gnse9FUgKD5fs= | Specifies the SSH key for the user account (usam). |
| | switch(config)# no username admin sshkey ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEAtjIHrIt/3dDeohix6JcRSIYZ 0EOdJ3l5RONWcwSgAuTUSrLk 3a9hdYkzY94fhHmNGQGCjVg+8cbOxyH4Z1jcVFcrDogtQT+Q8dve qts/8XQhqkNAFeGy4u8TJ2Us oreCU6DlibwkpzDafzKTPA5vB6FmHd2TI6Gnse9FUgKD5fs= | Deletes the SSH key for the user account (usam). |

Send documentation comments to mdsfeedback-doc@cisco.com.

Logging out Users

To log out another user on the switch, use the **clear user** command.

In the following example, the user named vsam is logged out from the switch.

```
switch# clear user vsam
```

Use the **show users** command to view a list of the logged in users (see [Example 19-17](#)).

Example 19-17 Displays All Logged in Users

```
switch# show users
admin pts/7 Jan 12 20:56 (10.77.202.149)
admin pts/9 Jan 12 23:29 (modena.cisco.com)
admin pts/10 Jan 13 03:05 (dhcp-171-71-58-120.cisco.com)
admin pts/11 Jan 13 01:53 (dhcp-171-71-49-49.cisco.com)
```

Displaying User Account Information

Use the **show user-account** command to display configured information about user accounts. See Examples [19-18](#) to [19-19](#).

Example 19-18 Displays Information for a Specified User

```
switch# show user-account user1
user:user1
    this user account has no expiry date
    roles:network-operator
no password set. Local login not allowed
Remote login through RADIUS is possible
```

Example 19-19 Displays Information for All Users

```
switch# show user-account
show user-account
user:admin
    this user account has no expiry date
    roles:network-admin
user:usam
    expires on Sat May 31 00:00:00 2003
    roles:network-admin network-operator
user:msam
    this user account has no expiry date
    roles:network-operator
user:user1
    this user account has no expiry date
    roles:network-operator
no password set. local login not allowed
Remote login through RADIUS is possible
```

Send documentation comments to mdsfeedback-doc@cisco.com.

SNMP Security

From Cisco MDS SAN-OS Release 1.2, the CLI and SNMP use common roles in all switches in the Cisco MDS 9000 Family. You can use SNMP to modify a role that was created using CLI and vice versa (see [Chapter 22, “Configuring SNMP”](#)).

From Cisco MDS SAN-OS Release 2.0(1b), users, passwords, and roles for all CLI and SNMP users are the same. A user configured through CLI can access the switch using the SNMP (for example, the Fabric Manager or the Device Manager) and vice versa.

Configuring Accounting Services

Accounting refers to the log information that is kept for each management session in a switch. This information may be used to generate reports for troubleshooting and auditing purposes. Accounting can be implemented locally or remotely (using RADIUS).



Tip

The Cisco MDS 9000 Family switch uses interim-update RADIUS accounting-request packets to communicate accounting log information to the RADIUS server. The RADIUS server must be appropriately configured to log the information communicated in these packets. Several servers typically have log update/watchdog packets flags in the AAA client configuration. Turn on this flag to ensure proper RADIUS accounting.



Note

Configuration operations are automatically recorded in the accounting log if they are performed in configuration mode. Additionally, important system events (for example, configuration save and system switchover) are also recorded in the accounting log.

Displaying Accounting Configuration

The **show accounting** command displays configured accounting information. See Examples [19-20](#) to [19-22](#). To specify the size of the local accounting log to be displayed, use the **show accounting log** command. By default about 250KB of accounting log is displayed.

Example 19-20 Displays Two Samples of Configured Accounting Parameters

```
switch# show accounting config
show aaa accounting
      default: local

switch# show aaa accounting
      default: group rad1
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 19-21 Displays 60K of the Accounting Log

```
switch# show accounting log 600000
Fri Jan 16 15:28:21 1981:stop:snmp_348506901_64.104.131.208:admin:
Fri Jan 16 21:17:04 1981:start:/dev/pts/0_348527824:admin:
Fri Jan 16 21:35:45 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group1
Fri Jan 16 21:35:51 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group1
Fri Jan 16 21:35:51 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group5
Fri Jan 16 21:35:55 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group5
Fri Jan 16 21:35:55 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group3
Fri Jan 16 21:58:17 1981:start:snmp_348530297_171.71.150.105:admin:
...
```

Example 19-22 Displays the Entire Log File.

```
switch# show accounting log
Fri Jan 16 15:28:21 1981:stop:snmp_348506901_64.104.131.208:admin:
Fri Jan 16 21:17:04 1981:start:/dev/pts/0_348527824:admin:
Fri Jan 16 21:35:45 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group1
Fri Jan 16 21:35:51 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group1
Fri Jan 16 21:35:51 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group5
Fri Jan 16 21:35:55 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group5
Fri Jan 16 21:35:55 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group3
Fri Jan 16 21:58:17 1981:start:snmp_348530297_171.71.150.105:admin:
Fri Jan 16 21:58:17 1981:stop:snmp_348530297_171.71.150.105:admin:
Fri Jan 16 21:58:18 1981:start:snmp_348530298_171.71.150.105:admin:
Fri Jan 16 21:58:18 1981:stop:snmp_348530298_171.71.150.105:admin:
...
Fri Jan 16 23:37:02 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group3
Fri Jan 16 23:37:26 1981:update:/dev/pts/0_348527824:admin:updated TACACS+ parameters for
group:TacacsServer1
Fri Jan 16 23:45:19 1981:update:/dev/pts/0_348527824:admin:updated TACACS+ parameters for
group:TacacsServer1
Fri Jan 16 23:45:19 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group1
...
Fri Jan 16 23:53:51 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
server:Server3
Fri Jan 16 23:54:00 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
server:Server5
Fri Jan 16 23:54:22 1981:update:/dev/pts/0_348527824:admin:updated TACACS+ parameters for
server:ServerA
Fri Jan 16 23:54:25 1981:update:/dev/pts/0_348527824:admin:updated TACACS+ parameters for
server:ServerB
Fri Jan 16 23:55:03 1981:update:/dev/pts/0_348527824:admin:updated RADIUS parameters for
group:Group1
...
Sat Jan 17 00:01:41 1981:start:snmp_348537701_171.71.58.100:admin:
Sat Jan 17 00:01:41 1981:stop:snmp_348537701_171.71.58.100:admin:
Sat Jan 17 00:01:42 1981:start:snmp_348537702_171.71.58.100:admin:
Sat Jan 17 00:01:42 1981:stop:snmp_348537702_171.71.58.100:admin:
...
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Clearing Accounting Logs

To clear out the contents of the current log, use the **clear accounting log** command.

```
switch# clear accounting log
```

Configuring SSH Services

The Telnet service is enabled by default on all Cisco MDS 9000 Family switches. Before enabling the SSH service, generate a server key pair. (see the “[Generating the SSH Server Key Pair](#)” section on [page 19-35](#)).

Use the **ssh key** command to generate a server key.

Enabling SSH Service

By default, the SSH service is disabled.

To enable or disable the SSH service, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# ssh server enable updated | Enables the use of the SSH service. |
| | switch(config)# no ssh server enable updated | Disables (default) the use of the SSH service and resets the switch to its factory defaults. |



Caution

If you are logging in to a switch through SSH and you have issued the **aaa authentication login default none** command, you must enter one or more key strokes to log in. If you press the **Enter** key without entering at least one keystroke, your log in will be rejected.

Specifying the SSH Key

You can specify a SSH key to log in using the SSH client without being prompted for a password.

To specify or delete the SSH Key for a specified user, follow these steps:

| | Command | Purpose |
|--------|-------------------------|----------------------------|
| Step 1 | switch# config t | Enters configuration mode. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|---|--|
| Step 2 | <pre>switch(config)# username admin sshkey ssh-rsa AAAAAB3NzaC1yc2EAAAABIwAAAIEAtjIHrIt/3dDeohix6JcRSIYZ 0EOdJ315RONWcwSgAutUSrLk3a9hdYkzY94fhHmNGQGCjVg+8cbO xyH4Z1jcVFcrDogtQT+Q8dveqts/8XQhqkNAFeGy4u8TJ2UsoreC U6D1libwkpzDafzKTPA5vB6FmHd2TI6Gnse9FUGKD5fs=</pre> | Specifies the SSH key for the user account (usam). |
| | <pre>switch(config)# no username admin sshkey ssh-rsa AAAAAB3NzaC1yc2EAAAABIwAAAIEAtjIHrIt/3dDeohix6JcRSIYZ 0EOdJ315RONWcwSgAutUSrLk3a9hdYkzY94fhHmNGQGCjVg+8cbO xyH4Z1jcVFcrDogtQT+Q8dveqts/8XQhqkNAFeGy4u8TJ2UsoreC U6D1libwkpzDafzKTPA5vB6FmHd2TI6Gnse9FUGKD5fs=</pre> | Deletes the SSH key for the user account (usam). |

Generating the SSH Server Key Pair

Be sure to have an SSH server key pair with the appropriate version before enabling the SSH service. Generate the SSH server key pair according to the SSH client version used. The number of bits specified for each key pair ranges from 768 to 2048.

The SSH service accepts three types of key pairs for use by SSH versions 1 and 2.

- The **rsa1** option generates the RSA1 key pair for the SSH version 1 protocol.
- The **dsa** option generates the DSA key pair for the SSH version 2 protocol.
- The **rsa** option generates the RSA key pair for the SSH version 2 protocol.



Caution If you delete all of the SSH keys, you cannot start a new SSH session.

To generate the SSH server key pair, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | <pre>switch(config)# ssh key rsa1 1024 generating rsa1 key..... generated rsa1 key</pre> | Generates the RSA1 server key pair. |
| | <pre>switch(config)# ssh key dsa 1024 generating dsa key..... generated dsa key</pre> | Generates the DSA server key pair. |
| | <pre>switch(config)# ssh key rsa 1024 generating rsa key..... generated rsa key</pre> | Generates the RSA server key pair. |
| | <pre>switch(config)# no ssh key rsa 1024 cleared RSA keys</pre> | Clears the RSA server key pair configuration. |

Overwriting a Generated Key Pair

If the SSH key pair option is already generated for the required version, use the **force** option to overwrite the previously generated key pair.

Send documentation comments to mdsfeedback-doc@cisco.com.

To overwrite the previously generated key pair, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# confi g t | Enters configuration mode. |
| Step 2 | switch(config)# ssh key dsa 768 ssh key dsa 512 dsa keys already present, use force option to overwrite them switch(config)# ssh key dsa 512 force deleting old dsa key..... generating dsa key..... generated dsa key | Tries to set the server key pair. If a required server key pair is already configured, use the force option to overwrite that server key pair. Deletes the old DSA key and sets the server key pair using the new bit specification. |

Clearing SSH Hosts

The **clear ssh hosts** command clears the existing list of trusted SSH hosts and reallows you to use SCP/SFTP along with **copy** command for particular hosts.

When you use SCP/SFTP along with the **copy** command, a list of trusted SSH hosts are built and stored within the switch (see [Example 19-23](#)).

Example 19-23 Using SCP/SFTP to Copy Files

```
switch# copy scp://abcd@171.71.48.223/users/abcd/abc
bootflash:abc The authenticity of host '171.71.48.223 (171.71.48.223)'
can't be established.
RSA1 key fingerprint is 01:29:62:16:33:ff:f7:dc:cc:af:aa:20:f8:20:a2:db.
Are you sure you want to continue connecting (yes/no)? yes
Added the host to the list of known hosts
(/var/home/admin/.ssh/known_hosts). [SSH key information about the host is
stored on the switch]
abcd@171.71.48.223's password:
switch#
```

If a host's SSH key changes before you use SCP/SFTP along with the **copy** command, you will receive an error (see [Example 19-24](#)).

Example 19-24 Using SCP/SFTP to Copy Files—Error Caused by SSH Key Change

```
switch# copy scp://apn@171.69.16.46/isan-104
bootflash:isan-ram-1.0.4
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
@  WARNING: REMOTE HOST IDENTIFICATION HAS CHANGED!  @
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
IT IS POSSIBLE THAT SOMEONE IS DOING SOMETHING NASTY!
Someone could be eavesdropping on you right now (man-in-the-middle attack)!
It is also possible that the RSA1 host key has just been changed.
The fingerprint for the RSA1 key sent by the remote host is
36:96:ca:d7:29:99:79:74:aa:4d:97:49:81:fb:23:2f.
Please contact your system administrator.
Add correct host key in /mnt/pss/.ssh/known_hosts to get rid of this
message.
Offending key in /mnt/pss/.ssh/known_hosts:2
RSA1 host key for 171.69.16.46 has changed and you have requested strict
checking.
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying SSH Protocol Status

Use the **show ssh server** command to display the status of the SSH protocol (enabled or disabled) and the versions that are enabled for that switch (see [Example 19-25](#)).

Example 19-25 Displays SSH Protocol Status

```
switch# show ssh server
ssh is enabled
version 1 enabled
version 2 enabled
```

Use the **show ssh key** command to display the server key pair details for the specified key or for all keys, (see [Example 19-26](#)).

Example 19-26 Displays Server Key Pair Details

```
switch# show ssh key
rsa1 Keys generated:Sun Jan 13 07:16:26 1980
1024 35
fingerprint:
1024 67:76:02:bd:3e:8d:f5:ad:59:5a:1e:c4:5e:44:03:07
could not retrieve rsa key information
dsa Keys generated:Sun Jan 13 07:40:08 1980
ssh-dss
AAAAB3NzaC1kc3MAAABBAJTCRQOydNRel2v7uiO6Fix+OTn8eGdnnDVxw5eJs5OcOEXOyjaWcMMYsEgxc9ada1NElp
8Wy7GPMWGOQYj9CU0AAAAVAMCcWhNN18zFNOIPo7cU3t7d0iEbAAAAQBdQ8UAOi/Cti84qFb3kTqXlS9mEhdQUo0lH
cH5bw5PKfj2Y/dLR437zCBKXetPj4p7mhQ6Fq5os8RZtJEyOsNsAAABAA0oxZbPyWeR5NHATXiyXdPI7j9i8fgyn9F
NipMkOF2Mn75Mi/lqQ4NIq0gQNvQ0x27uCeQlRts/QwI4q68/eaw=
fingerprint:
512 f7:cc:90:3d:f5:8a:a9:ca:48:76:9f:f8:6e:71:d4:ae
```

Recovering Administrator Password

An administrator can recover a password from a local console connection. The password recovery procedure must be performed on the supervisor module that becomes the active supervisor module after the recovery procedure is completed.

To ensure the other supervisor module does not become the active module, you have two options:

- Physically remove the other supervisor module from the chassis, or
- For the duration of this procedure, change the other supervisor module's console prompt to the loader> or switch(boot) # prompt (see [Chapter 6, “Software Images”](#)).



Note

Password recovery is not possible from a Telnet or SSH session.

To recover a administrator's password, follow these steps:

Step 1 Reboot the switch.

```
switch# reload
The supervisor is going down for reboot NOW!
```

Send documentation comments to mdsfeedback-doc@cisco.com.

- Step 2** Press the **Ctrl-]** key sequence (when the switch begins its Cisco SAN-OS software boot sequence) to enter the `switch(boot)#` prompt (see the “Recovery from the `switch(boot)#` Prompt” section on [page 6-31](#)).

```
Ctrl-]
switch(boot)#
```

- Step 3** Change to configuration mode.

```
switch(boot)# config terminal
```

- Step 4** Enter the **admin-password** command to reset the administrator password.

```
switch(boot-config)# admin-password password
```

- Step 5** Exit to the EXEC mode.

```
switch(boot-config)# exit
switchboot#
```

- Step 6** Enter the **load** command to load the Cisco SAN-OS software.

```
switch(boot)# load bootflash:system.img
```

- Step 7** Save the software configuration.

```
switch# copy running-config startup-config
```

Configuring Cisco ACS Servers

The Cisco Access Control Server (ACS) uses TACACS+ and RADIUS protocols to provide AAA services that ensure a secure environment. When using the AAA server, user management is normally done using Cisco ACS. [Figure 19-2](#), [Figure 19-3](#), [Figure 19-4](#), and [Figure 19-5](#) display ACS server user setup configurations for network-admin roles and multiple roles using either TACACS+ or RADIUS.

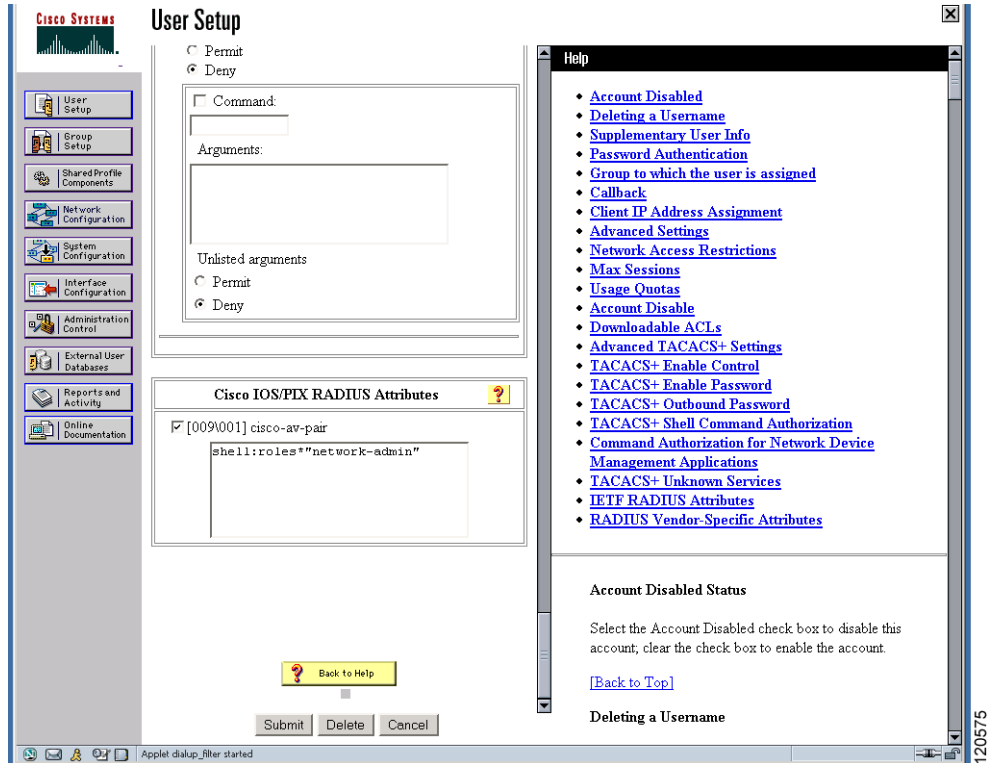


Caution

Cisco MDS SAN-OS does not support all numeric usernames, whether created with TACACS+ or RADIUS, or created locally. Local users with all numeric names cannot be created. If an all numeric username exists on an AAA server and is entered during login, the user is not logged in.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 19-2 Configuring the network-admin Role When Using RADIUS



Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 19-3 Configuring Multiple Roles with SNMPv3 Attributes When Using RADIUS

The screenshot shows the CiscoSecure ACS web interface for User Setup. The browser address bar shows `http://10.76.100.108:2691/index2.htm`. The interface includes a left sidebar with navigation links: User Setup, Group Setup, Shared Profile Components, Network Configuration, System Configuration, Interface Configuration, Administration Control, External User Databases, Reports and Activity, and Online Documentation. The main content area is titled "User Setup" and contains the following sections:

- Per User Command Authorization:**
 - Unmatched Cisco IOS commands:
 - ☐ Permit
 - ☒ Deny
 - Command:
 - Arguments:
 - Unlisted arguments:
 - ☐ Permit
 - ☒ Deny
- Cisco IOS/PIX RADIUS Attributes:**
 - ☒ [009A001] cisco-av-pair


```
shell:roles="Role1 Role3 Role5
Role7"snmpv3:auth=MDS priv=DES
```

At the bottom of the main content area are buttons for **Submit**, **Delete**, and **Cancel**.

On the right side, there is a **Help** pane with a list of links:

- [Account Disabled](#)
- [Deleting a Username](#)
- [Supplementary User Info](#)
- [Password Authentication](#)
- [Group to which the user is assigned](#)
- [Callback](#)
- [Client IP Address Assignment](#)
- [Advanced Settings](#)
- [Network Access Restrictions](#)
- [Max Sessions](#)
- [Usage Quotas](#)
- [Account Disable](#)
- [Downloadable ACLs](#)
- [Advanced TACACS+ Settings](#)
- [TACACS+ Enable Control](#)
- [TACACS+ Enable Password](#)
- [TACACS+ Outbound Password](#)
- [TACACS+ Shell Command Authorization](#)
- [Command Authorization for Network Device Management Applications](#)
- [TACACS+ Unknown Services](#)
- [IETF RADIUS Attributes](#)
- [RADIUS Vendor-Specific Attributes](#)

Below the links, there is a section titled **Account Disabled Status** with the text: "Select the Account Disabled check box to disable this account; clear the check box to enable the account." and a [\[Back to Top\]](#) link.

At the bottom of the help pane, there is a section titled **Deleting a Username**.

The status bar at the bottom of the browser window shows "Applet dialup_filter started" and the page number "120576".

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 19-4 Configuring the network-admin Role with SNMPv3 Attributes When Using TACACS+

The screenshot shows the 'User Setup' page in the Cisco Systems web interface. The left sidebar contains navigation links: User Setup, Group Setup, Shared Profile Components, Network Configuration, System Configuration, Interface Configuration, Administration Control, External User Databases, Reports and Activity, and Online Documentation. The main content area is titled 'TACACS+ Settings' and includes a 'Help' button. The settings are organized into two sections: 'PPP IP' and 'Shell (exec)'. The 'PPP IP' section has checkboxes for 'In access control list', 'Out access control list', 'Route', 'Routing', and 'Custom attributes'. The 'Routing' checkbox is checked, and the 'Enabled' checkbox is also checked. The 'Shell (exec)' section has checkboxes for 'Access control list', 'Auto command', 'Callback line', 'Callback rotary', 'Idle time', 'No callback verify', 'No escape', 'No hangup', 'Privilege level', 'Timeout', and 'Custom attributes'. The 'Custom attributes' checkbox is checked. Below the checkboxes, there is a text area containing the command: `cisco-av-pair=shell:roles="Role1 Role3"snmpv3:auth=MDS priv=DES`. At the bottom of the settings area are 'Submit', 'Delete', and 'Cancel' buttons. On the right side, there is a 'Help' panel with a list of links: Account Disabled, Deleting a Username, Supplementary User Info, Password Authentication, Group to which the user is assigned, Callback, Client IP Address Assignment, Advanced Settings, Network Access Restrictions, Max Sessions, Usage Quotas, Account Disable, Downloadable ACLs, Advanced TACACS+ Settings, TACACS+ Enable Control, TACACS+ Enable Password, TACACS+ Outbound Password, TACACS+ Shell Command Authorization, Command Authorization for Network Device Management Applications, TACACS+ Unknown Services, IETF RADIUS Attributes, and RADIUS Vendor-Specific Attributes. Below the links, there is a section titled 'Account Disabled Status' with instructions: 'Select the Account Disabled check box to disable this account; clear the check box to enable the account.' and a '[Back to Top]' link. Below that is a section titled 'Deleting a Username' with instructions: 'The Delete button appears only when you are editing an...'. The status bar at the bottom shows 'Applet dialup_filter started'.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 19-5 Configuring Multiple Roles with SNMPv3 Attributes When Using TACACS+

User Setup

TACACS+ Settings

☐ PPP IP

☐ In access control list

☐ Out access control list

☐ Route

☐ Routing

☐ Custom attributes

☐ Enabled

Note: PPP LCP will be automatically enabled if this service is enabled

☒ **Shell (exec)**

☐ Access control list

☐ Auto command

☐ Callback line

☐ Callback rotary

☐ Idle time

☐ No callback verify

☐ No escape

☐ No hangup

☐ Privilege level

☐ Timeout

☒ Custom attributes

cisco-av-pair*shell:roles=
network-admin*snmpv3:auth=md5
priv=aes-128

Submit Delete Cancel

Help

- [Account Disabled](#)
- [Deleting a Username](#)
- [Supplementary User Info](#)
- [Password Authentication](#)
- [Group to which the user is assigned](#)
- [Callback](#)
- [Client IP Address Assignment](#)
- [Advanced Settings](#)
- [Network Access Restrictions](#)
- [Max Sessions](#)
- [Usage Quotas](#)
- [Account Disable](#)
- [Downloadable ACLs](#)
- [Advanced TACACS+ Settings](#)
- [TACACS+ Enable Control](#)
- [TACACS+ Enable Password](#)
- [TACACS+ Outbound Password](#)
- [TACACS+ Shell Command Authorization](#)
- [Command Authorization for Network Device Management Applications](#)
- [TACACS+ Unknown Services](#)
- [IETF RADIUS Attributes](#)
- [RADIUS Vendor-Specific Attributes](#)

Account Disabled Status

Select the Account Disabled check box to disable this account; clear the check box to enable the account.

[\[Back to Top\]](#)

Deleting a Username

The Delete button appears only when you are editing

Default Settings

Table 19-2 lists the default settings for all switch security features in any switch.

Table 19-2 Default Switch Security Settings

| Parameters | Default |
|-----------------------------|--------------------------------------|
| Roles in Cisco MDS Switches | Network operator (network-operator). |
| AAA configuration services | Local. |
| Authentication port | 1821. |
| Accounting port | 1813. |
| Preshared key communication | Clear text. |
| RADIUS server time out | 1 (one) second. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 19-2 **Default Switch Security Settings (continued)**

| Parameters | Default |
|-------------------------|--------------------------------|
| RADIUS server retries | Once. |
| TACACS | Disabled. |
| TACACS servers | None configured. |
| TACACS server timeout | 5 seconds. |
| AAA server distribution | Disabled. |
| VSAN policy for roles | Permit. |
| User account | No expiry (unless configured). |
| Password | None. |
| Accounting log size | 250 KB. |
| SSH service | Disabled. |
| Telnet service | Enabled. |

Send documentation comments to mdsfeedback-doc@cisco.com.



Configuring Fabric Security

Fibre Channel Security Protocol (FC-SP) capabilities in Cisco MDS SAN-OS Release 1.3 provides switch-switch and host-switch authentication to overcome security challenges for enterprise-wide fabrics. Diffie-Hellman Challenge Handshake Authentication Protocol (DHCHAP) is an FC-SP protocol implemented in this release to provide authentication between Cisco MDS 9000 Family switches and other devices. It consists of the CHAP protocol combined with the Diffie-Hellman exchange.

This chapter includes the following sections:


- [About Fabric Authentication, page 20-2](#)
- [About DHCHAP, page 20-3](#)
- [DHCHAP Compatibility with Existing Cisco MDS Features, page 20-3](#)
- [Configuring DHCHAP Authentication, page 20-3](#)
- [DHCHAP Configuration, page 20-4](#)
- [DHCHAP Authentication Modes, page 20-4](#)
- [DHCHAP Hash Algorithm Configuration, page 20-5](#)
- [DHCHAP Group Configuration, page 20-6](#)
- [DHCHAP Password Configuration, page 20-6](#)
- [Password Configuration for Other Devices, page 20-7](#)
- [DHCHAP Timeout Value, page 20-8](#)
- [Displaying Protocol Security Information, page 20-9](#)
- [DHCHAP AAA Authentication, page 20-10](#)
- [Sample Configuration, page 20-10](#)
- [Default Settings, page 20-12](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

About Fabric Authentication

All switches in the Cisco MDS 9000 Family enable fabric-wide authentication from one switch to another switch, or from a switch to a host. These switch and host authentications are performed locally or remotely in each fabric. As storage islands are consolidated and migrated to enterprise-wide fabrics new security challenges arise. The approach of securing storage islands cannot always be guaranteed in enterprise-wide fabrics. For example, in a campus environment with geographically distributed switches someone could maliciously interconnect incompatible switches or you could accidentally do so, resulting in Inter-Switch Link (ISL) isolation and link disruption. This need for physical security is addressed by switches in the Cisco MDS 9000 Family (see [Figure 20-1](#)).

Figure 20-1 Switch and Host Authentication



105209



Note

Fibre Channel (FC) host bus adapters (HBAs) with appropriate firmware and drivers are required for host-switch authentication.

Send documentation comments to mdsfeedback-doc@cisco.com.

About DHCHAP

DHCHAP is an authentication protocol that authenticates the devices connecting to a switch. Fibre Channel authentication allows only trusted devices to be added to a fabric, thus preventing unauthorized devices from accessing the switch.

**Note**

The terms FC-SP and DHCHAP are used interchangeably in this chapter.

DHCHAP is a mandatory password-based, key-exchange authentication protocol that supports both switch-to-switch and host-to-switch authentication. DHCHAP negotiates hash algorithms and DH groups before performing authentication. It supports MD5 and SHA-1 algorithm-based authentication.

Configuring the DHCHAP feature requires the ENTERPRISE_PKG license (see [Chapter 3, “Obtaining and Installing Licenses”](#)).

DHCHAP Compatibility with Existing Cisco MDS Features

This section identifies the impact of configuring the DHCHAP feature along with existing Cisco MDS features:

- PortChannel interfaces—If DHCHAP is enabled for ports belonging to a PortChannel, DHCHAP authentication is performed at the physical interface level, not at the PortChannel level.
- FCIP interfaces—The DHCHAP protocol works with the FCIP interface just as it would with a physical interface.
- Port security or fabric binding—Fabric binding policies are enforced based on identities authenticated by DHCHAP.
- VSANs—DHCHAP authentication is not done on a per-VSAN basis.
- High availability—DHCHAP authentication works transparently with existing HA features.

Configuring DHCHAP Authentication

To configure DHCHAP authentication using the local password database, follow these steps:

-
- | | |
|---------------|--|
| Step 1 | Enable DHCHAP. |
| Step 2 | Identify and configure the DHCHAP authentication modes. |
| Step 3 | Configure the hash algorithm and DH group. |
| Step 4 | Configure the DHCHAP password for the local switch and other switches in the fabric. |
| Step 5 | Configure the DHCHAP timeout value for reauthentication. |
| Step 6 | Verify the DHCHAP configuration. |
-

Send documentation comments to mdsfeedback-doc@cisco.com.

DHCHAP Configuration

By default, the DHCHAP feature is disabled in all switches in the Cisco MDS 9000 Family. You must explicitly enable the DHCHAP feature to access the configuration and verification commands for fabric authentication. When you disable this feature, all related configurations are automatically discarded.

To enable DHCHAP for a Cisco MDS switch, follow these steps:

| | Command | Purpose |
|--------|---------------------------------------|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# fcsp enable | Enables the DHCHAP in this switch. |
| | switch(config)# no fcsp enable | Disables (default) the DHCHAP in this switch. |

DHCHAP Authentication Modes

The DHCHAP authentication status for each interface depends on the configured DHCHAP port mode. When the DHCHAP feature is enabled in a switch, each Fibre Channel interface or FCIP interface may be configured to be in one of four DHCHAP port modes:

- On—During switch initialization, if the connecting device supports DHCHAP authentication, the software performs the authentication sequence. If the connecting device does not support DHCHAP authentication, the software moves the link to an isolated state.
- Auto-Active—During switch initialization, if the connecting device supports DHCHAP authentication, the software performs the authentication sequence. If the connecting device does not support DHCHAP authentication, the software continues with the rest of the initialization sequence.
- Auto-Passive (default)—The switch does not initiate DHCHAP authentication, but participates in DHCHAP authentication if the connecting device initiates DHCHAP authentication.
- Off—The switch does not support DHCHAP authentication. Authentication messages sent to such ports return error messages to the initiating switch.



Note

Whenever DHCHAP port mode is changed to a mode other than the Off mode, reauthentication is performed.

Table 20-1 identifies the switch-to-switch authentication behavior between two Cisco MDS switches in various modes.

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 20-1 *DHCHAP Authentication Status Between Two MDS Switches*

| Switch N DHCHAP Modes | Switch 1 DHCHAP Modes | | | |
|-----------------------------|------------------------------------|---|---|---|
| | on | auto-active | auto-passive | off |
| on | FC-SP authentication is performed. | FC-SP authentication is performed. | FC-SP authentication is performed. | Link is brought down. |
| auto-Active | | | | |
| auto-Passive | | | FC-SP authentication is <i>not</i> performed. | FC-SP authentication is <i>not</i> performed. |
| off | Link is brought down. | FC-SP authentication is <i>not</i> performed. | | |

To enable the DHCHAP mode for a particular interface, follow these steps:

| | Command | Purpose |
|---------------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc2/1-3 switch(config-if)# | Enters the interface submode. |
| Step 3 | switch(config-if)# fcsp on | Sets the DHCHAP mode for the selected interfaces to be in the on state. |
| | switch(config-if)# no fcsp on | Reverts to the factory default of auto-passive for these three interfaces. |
| Step 4 | switch(config-if)# fcsp auto-active 0 | Changes the DHCHAP authentication mode for the selected interfaces to auto-active. The 0 indicates that the port does not perform authentication |
| | switch(config-if)# fcsp auto-active 120 | Changes the DHCHAP authentication mode to auto-active for the selected ports to reauthenticate every two hours (120 minutes) after the initialization authentication. |
| | switch(config-if)# fcsp auto-active | Changes the DHCHAP authentication mode to auto-active for the selected ports. |

DHCHAP Hash Algorithm Configuration

Cisco MDS switches support a default hash algorithm priority list of MD5 followed by SHA-1 for DHCHAP authentication.



Tip

If you change the hash algorithm configuration, then change it globally for all switches in the fabric.



Caution

RADIUS and TACACS+ protocols always use MD5 for CHAP authentication. Using SHA-1 as the hash algorithm may prevent RADIUS and TACACS+ usage—even if these AAA protocols are enabled for DHCHAP authentication.

Send documentation comments to mdsfeedback-doc@cisco.com.

To change the hash algorithm, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# fcsp dhchap hash sha1 | Configures the use of only the SHA-1 hash algorithm. |
| | switch(config)# fcsp dhchap hash MD5 | Configures the use of only the MD5 hash algorithm. |
| | switch(config)# fcsp dhchap hash md5 sha1 | Defines the use of the default hash algorithm priority list of MD5 followed by SHA-1 for DHCHAP authentication. |
| | switch(config)# no fcsp dhchap hash sha1 | Reverts to the factory default priority list of the MD5 hash algorithm followed by the SHA-1 hash algorithm. |

DHCHAP Group Configuration

All switches in the Cisco MDS Family support all DHCHAP groups specified in the standard: 0 (null DH group, which does not perform the Diffie-Hellman exchange), 1, 2, 3, or 4.



Tip

If you change the DH group configuration, change it globally for all switches in the fabric.

To change the DH group settings, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# fcsp dhchap group 2 3 4 | Prioritizes the use of DH group 2, 3, and 4 in the configured order. |
| | switch(config)# no fcsp dhchap group 0 | Reverts to the DHCHAP factory default order of 0, 4, 1, 2, and 3. |

DHCHAP Password Configuration

DHCHAP authentication in each direction requires a shared secret password between the connected devices. To do this, you can use one of three approaches to manage passwords for all switches in the fabric that participate in DHCHAP.

- Approach 1—Use the same password for all switches in the fabric. This is the simplest approach. When you add a new switch, you use the same password to authenticate that switch in this fabric. It is also the most vulnerable approach if someone from the outside maliciously attempts to access any one switch in the fabric
- Approach 2—Use a different password for each switch and maintain that password list in each switch in the fabric. When you add a new switch, you create a new password list and update all switches with the new list. Accessing one switch yields the password list for all switches in that fabric.
- Approach 3—Use different passwords for different switches in the fabric. When you add a new switch, multiple new passwords corresponding to each switch in the fabric must be generated and configured in each switch. Even if one switch is compromised, the password of other switches are still protected. This approach requires considerable password maintenance by the user.

Send documentation comments to mdsfeedback-doc@cisco.com.

**Note**

All passwords are restricted to 64 alphanumeric characters and can be changed, but not deleted.

**Tip**

We recommend using RADIUS or TACACS+ for fabrics with more than five switches. If you need to use a local password database, you can continue to do so using Approach 3 and using the Cisco MDS 9000 Family Fabric Manager to manage the password database.

Refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide* for further information.

Configuring the DHCHAP Password for the Local Switch

To configure the DHCHAP password for the local switch, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | <code>switch# config t</code> | Enters configuration mode. |
| Step 2 | <code>switch(config)# fcsp dhchap password 0 mypassword</code> | Configures a clear text password for the local switch. |
| | <code>switch(config)# fcsp dhchap password 0 mypassword 30:11:bb:cc:dd:33:11:22</code> | Configures a clear text password for the local switch to be used for the device with the specified WWN. |
| | <code>switch(config)# no fcsp dhchap password 0 mypassword 30:11:bb:cc:dd:33:11:22</code> | Removes the clear text password for the local switch to be used for the device with the specified WWN. |
| | <code>switch(config)# fcsp dhchap password 7 sfsfdf</code> | Configures a password entered in an encrypted format for the local switch. |
| | <code>switch(config)# fcsp dhchap password 7 sfsfdf 29:11:bb:cc:dd:33:11:22</code> | Configures a password entered in an encrypted format for the local switch to be used for the device with the specified WWN. |
| | <code>switch(config)# no fcsp dhchap password 7 sfsfdf 29:11:bb:cc:dd:33:11:22</code> | Removes the password entered in an encrypted format for the local switch to be used for the device with the specified WWN. |
| | <code>switch(config)# fcsp dhchap password mypassword1</code> | Configures a clear text password for the local switch to be used with any connecting device. |

Password Configuration for Other Devices

You can configure passwords in the local authentication database for other devices in a fabric. The other devices are identified by their device name, which is also known as the switch WWN or device WWN. The password is restricted to 64 characters and can be specified in clear text (0) or in encrypted text (7).

**Note**

The switch WWN identifies the physical switch. This WWN is used to authenticate the switch and is different from the VSAN node WWN.

Send documentation comments to mdsfeedback-doc@cisco.com.

Locally Configuring the Device Name

To locally configure the device name of another switch in the fabric, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# fcsp dhchap devicename 00:11:22:33:44:aa:bb:cc password NewPassword | Configures a password for another switch in the fabric that is identified by the switch WWN device name. |
| | switch(config)# no fcsp dhchap devicename 00:11:22:33:44:aa:bb:cc password NewPassword | Removes the password entry for this switch from the local authentication database. |
| | switch(config)# fcsp dhchap devicename 00:11:55:66:00:aa:bb:cc password 0 NewPassword | Configures a clear text password for another switch in the fabric that is identified by the switch WWN device name. |
| | switch(config)# fcsp dhchap devicename 00:11:22:33:55:aa:bb:cc password 7 asdf1kjh | Configures a password entered in an encrypted format for another switch in the fabric that is identified by the switch WWN device name. |

DHCHAP Timeout Value

During the DHCHAP protocol exchange, if the MDS switch does not receive the expected DHCHAP message within a specified time interval, authentication failure is assumed. The time ranges from 20 (no authentication is performed) to 1000 seconds. The default is 30 seconds.

When changing the timeout value, consider the following factors:

- The existing RADIUS and TACACS+ timeout values.
- The same value must also be configured all switches in the fabric.

Configuring the Timeout Value

To configure the DHCHAP timeout value, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# fcsp timeout 60 | Configures the reauthentication timeout to be 60 seconds. |
| | switch(config)# no fcsp timeout 60 | Reverts to the factory default of 30 seconds. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying Protocol Security Information

Use the **show fcsp** commands to display configurations for the local database (see [Example 20-1](#) through [20-6](#)).

Example 20-1 Displays DHCHAP Configurations in FC Interfaces

```
switch# show fcsp interface fc1/9

fc1/9:
  fcsp authentication mode:SEC_MODE_ON
  Status: Successfully authenticated
```

Example 20-2 Displays DHCHAP Statistics for a FC Interface

```
switch# show fcsp interface fc1/9 statistics

fc1/9:
  fcsp authentication mode:SEC_MODE_ON
  Status: Successfully authenticated
  Statistics:
  FC-SP Authentication Succeeded:5
  FC-SP Authentication Failed:0
  FC-SP Authentication Bypassed:0
```

Example 20-3 Displays the FC-SP WWN of the Device Connected through a Specified Interface

```
switch# show fcsp interface fc 2/1 wwn

fc2/1:
  fcsp authentication mode:SEC_MODE_ON
  Status: Successfully authenticated
  Other device's WWN:20:00:00:e0:8b:0a:5d:e7
```

Example 20-4 Displays Hash Algorithm and DHCHAP Groups Configured for the Local Switch

```
switch# show fcsp dhchap
Supported Hash algorithms (in order of preference):
DHCHAP_HASH_MD5
DHCHAP_HASH_SHA_1

Supported Diffie Hellman group ids (in order of preference):
DHCHAP_GROUP_NULL
DHCHAP_GROUP_1536
DHCHAP_GROUP_1024
DHCHAP_GROUP_1280
DHCHAP_GROUP_2048
```

Example 20-5 Displays the DHCHAP Local Password Database

```
switch# show fcsp dhchap database
DHCHAP Local Password:
  Non-device specific password:mypassword1
  Password for device with WWN:29:11:bb:cc:dd:33:11:22 is pjoalf
  Password for device with WWN:30:11:bb:cc:dd:33:11:22 is mypassword

Other Devices' Passwords:
  Password for device with WWN:00:11:22:33:44:aa:bb:cc is NewPassword
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 20-6 Displays the ASCII Representation of the Device WWN

```
switch# show fcsp asciiwnn 30:11:bb:cc:dd:33:11:22
Ascii representation of WWN to be used with AAA servers:Ox_3011bbccdd331122
```



Tip

Use the ASCII representation of the device WWN (identified in bold in [Example 20-6](#)) to configure the switch information on RADIUS and TACACS+ servers.

DHCHAP AAA Authentication

You can individually set authentication options. If authentication is not configured, local authentication is used by default.

Use the **aaa authentication dhchap** command to set authentication in any Cisco MDS 9000 Family switch.


To configure the AAA authentication, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# aaa authentication dhchap default group TacacsServer1 | Enables DHCHAP to use the TACACS+ server group (in this example, TacacsServer1) for authentication. |
| | switch(config)# aaa authentication dhchap default local | Enables DHCHAP for local authentication. |
| | switch(config)# aaa authentication dhchap default group RadiusServer1 | Enables DHCHAP to use the RADIUS server group (in this example, RadiusServer1) for authentication. |

Sample Configuration

This section provides the steps to configure the example illustrated in [Figure 20-2](#).

Figure 20-2 Sample DHCHAP Authentication



Send documentation comments to mdsfeedback-doc@cisco.com.

To configure the authentication setup shown in [Figure 20-2](#), follow these steps:

- Step 1** Obtain the device name of the MDS 9216 Switch in the fabric, The MDS 9216 Switch in the fabric is identified by the switch WWN.

```
MDS-9216# show wwn switch
Switch WWN is 20:00:00:05:30:00:54:de
```

- Step 2** Explicitly enable DHCHAP in this switch.



Note When you disable DHCHAP, all related configurations are automatically discarded.

```
MDS-9216(config)# fcsp enable
```

- Step 3** Configure a clear text password for this switch. This password will be used by the connecting device.

```
MDS-9216(config)# fcsp dhchap password rtp9216
```

- Step 4** Configures a password for another switch in the fabric that is identified by the switch WWN device name.

```
MDS-9216(config)# fcsp dhchap devicename 20:00:00:05:30:00:38:5e password rtp9509
```

- Step 5** Enable the DHCHAP mode for the required Fibre Channel interface.



Note Whenever DHCHAP port mode is changed to a mode other than the Off mode, reauthentication is performed.

```
MDS-9216(config)# interface fc 1/16
MDS-9216(config-if)# fcsp on
```

- Step 6** Verify the protocol security information configured in this switch by displaying the DHCHAP local password database.

```
MDS-9216# show fcsp dhchap database
DHCHAP Local Password:
    Non-device specific password:upt9216
Other Devices' Passwords:
    Password for device with WWN:20:00:00:05:30:00:38:5e is upt9509
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Step 7 Display the DHCHAP configuration in the Fibre Channel Interface

```
MDS-9216# show fcsp interface fc 1/6
fc1/6
      fcsp authentication mode:SEC_MODE_ON
      Status:Successfully authenticated
```

Step 8 Repeat these steps on the connecting MDS 9509 Switch.

```
MDS-9509# show wwn switch
Switch WWN is 20:00:00:05:30:00:38:5e
MDS-9509(config)# fcsp enable
MDS-9509(config)# fcsp dhchap password rtp9509
MDS-9509(config)# fcsp dhchap devicename 20:00:00:05:30:00:54:de password rtp9216
MDS-9509(config)# interface fc 4/5
MDS-9509(config-if)# fcsp on
MDS-9509# show fcsp dhchap database
DHCHAP Local Password:
      Non-device specific password:upt9509
Other Devices' Passwords:
      Password for device with WWN:20:00:00:05:30:00:54:de is upt9216
MDS-9509# show fcsp interface fc 4/5
Fc4/5
      fcsp authentication mode:SEC_MODE_ON
      Status:Successfully authenticated
```

You have now enabled and configured DHCHAP authentication for the sample setup in Figure

Default Settings

Table 20-2 lists the default settings for all fabric security features in any switch.

Table 20-2 **Default Fabric Security Settings**

| Parameters | Default |
|--|---|
| DHCHAP feature | Disabled. |
| DHCHAP hash algorithm | A priority list of MD5 followed by SHA-1 for DHCHAP authentication. |
| DHCHAP authentication mode | Auto-passive. |
| DHCHAP group default priority exchange order | 0, 4, 1, 2, and 3 respectively. |
| DHCHAP timeout value | 30 seconds. |



Configuring Port Security

All switches in the Cisco MDS 9000 Family provide port security features that reject intrusion attempts and report these intrusions to the administrator.



Note

Port security is only supported for Fibre Channel ports.

This chapter includes the following sections:

- [Port Security Features, page 21-2](#)
- [Port Security Initiation, page 21-2](#)
- [Port Security Manual Configuration, page 21-3](#)
- [Port Security Activation, page 21-4](#)
- [About AutoLearning, page 21-7](#)
- [Port Security Configuration Distribution, page 21-9](#)
- [Database Merge Guidelines, page 21-12](#)
- [Database Interaction, page 21-12](#)
- [Port Security Database Copy, page 21-13](#)
- [Port Security Database Deletion, page 21-14](#)
- [Port Security Database Cleanup, page 21-14](#)
- [Displaying Port Security Configurations, page 21-15](#)
- [Default Settings, page 21-18](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

Port Security Features

Typically, any Fibre Channel device in a SAN can attach to any SAN switch port and access SAN services based on zone membership. Port security features prevent unauthorized access to a switch port in the Cisco MDS 9000 Family:

- Login requests from unauthorized Fibre Channel devices (Nx ports) and switches (xE ports) are rejected.
- All intrusion attempts are reported to the SAN administrator through system messages.
- As of Cisco SAN-OS Release 2.0(1b), configuration distribution happens using the CFS infrastructure. and is limited to those switches that are CFS capable. Distribution is disabled by default.
- Configuring the port security policy requires the ENTERPRISE_PKG license (see [Chapter 3, “Obtaining and Installing Licenses”](#)).

Port Security Enforcement

To enforce port security, configure the devices and switch port interfaces through which each device or switch is connected, and activate the configuration.

- Use the port world wide name (pWWN) or the node world wide name (nWWN) to specify the Nx port connection for each device.
- Use the switch world wide name (sWWN) to specify the xE port connection for each switch.

Each Nx and xE port can be configured to restrict a single port or a range of ports.

Enforcement of port security policies are done on every activation and when the port tries to come up.

The port security feature uses two databases to accept and implement configurations.

- Configuration database—All configuration changes are stored in the configuration database.
- Active database—The database currently enforced by the fabric. The port security feature requires all devices connecting to a switch to be part of the port security active database. The software uses this active database to enforce authorization.

Port Security Initiation

By default, the port security feature is disabled in all switches in the Cisco MDS 9000 Family.

To enable port security, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# port-security enable | Enables port security on that switch. |
| | switch(config)# no port-security enable | Disables (default) port security on that switch. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Port Security Manual Configuration

To configure port security on any switch in the Cisco MDS 9000 Family, follow these steps:

-
- | | |
|---------------|--|
| Step 1 | Identify the WWN of the ports that need to be secured. |
| Step 2 | Secure the fWWN to an authorized nWWN or pWWN. |
| Step 3 | Activate the port security database. |
| Step 4 | Verify your configuration. |
-

WWN Identification

If you decide to manually configure port security, be sure to adhere to the following guidelines:

- Identify switch ports by the interface or by the fWWN.
- Identify devices by the pWWN or by the nWWN.
- If an Nx port:
 - is allowed to login to SAN switch port Fx, then that Nx port can only log in through the specified Fx port.
 - nWWN is bound to a Fx port WWN, then all pWWNs in the Nx port are implicitly paired with the Fx port.
- TE port checking is done on each VSAN in the allowed VSAN list of the trunk port.
- All PortChannel xE ports must be configured with the same set of WWNs in the same PortChannel.
- E port security is implemented in the port VSAN of the E port. In this case the sWWN is used to secure authorization checks.
- Once activated, the config database can be modified without any effect on the active database.
- By saving the running configuration, you save the configuration database and activated entries in the active database. Learned entries in the active database are not saved.

Authorized Port Pair Addition

After identifying the WWN pairs that need to be bound, add those pairs to the port security database.



Tip

As of Cisco SAN-OS Release 2.0(1b), the remote switch binding can be specified at the local switch. To specify the remote interfaces, you can use either the fWWN or sWWN-interface combination.

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure port security, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# confi g t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# port-security database vsan 1 switch(config-port-security)# | Enters the port security database mode for the specified VSAN. |
| | switch(config)# no port-security database vsan 1 switch(config)# | Deletes the port security configuration database from the specified VSAN. |
| Step 3 | switch(config-port-security)# swwn 20:01:33:11:00:2a:4a:66 interface port-channel 5 | Configures the specified sWWN to only log in through PortChannel 5. |
| | switch(config-port-security)# any-wwn interface fc1/1 - fc1/8 | Configures any WWN to log in through the specified interfaces. |
| | switch(config-port-security)# pwwn 20:11:00:33:11:00:2a:4a fwwn 20:81:00:44:22:00:4a:9e | Configures the specified pWWN to only log in through the specified fWWN. |
| | switch(config-port-security)# no pwwn 20:11:00:33:11:00:2a:4a fwwn 20:81:00:44:22:00:4a:9e | Deletes the specified pWWN configured in the previous step. |
| | switch(config-port-security)# nwwn 26:33:22:00:55:05:3d:4c fwwn 20:81:00:44:22:00:4a:9e | Configures the specified nWWN to log in through the specified fWWN. |
| | switch(config-port-security)# pwwn 20:11:33:11:00:2a:4a:66 | Configures the specified pWWN to log in through any port in the fabric. |
| | switch(config-port-security)# pwwn 20:11:33:11:00:2a:4a:66 swwn 20:00:00:0c:85:90:3e:80 | Configures the specified pWWN to log in through any interface in the specified switch. |
| | switch(config-port-security)# pwwn 20:11:33:11:00:2a:4a:66 swwn 20:00:00:0c:85:90:3e:80 interface fc3/1 | Configures the specified pWWN to log in through the specified interface in the specified switch. |
| | switch(config-port-security)# any-wwn interface fc3/1 | Configures any WWN to log in through the specified interface in any switch. |
| | switch(config-port-security)# no any-wwn interface fc2/1 | Deletes the wildcard configured in the previous step. |

Port Security Activation

By default, the port security feature is not activated in any switch in the Cisco MDS 9000 Family.

When you activate the port security feature, the following apply:

- Auto-learning is also automatically enabled. When auto-learning is enabled, the following apply:
 - From this point, learning happens only for the devices or interfaces that were not activated.
 - You will not be allowed to activate the database till you disable learning.
- All the logged-in devices are learned and are added to the active database
- All entries in the configured database are copied to the active database.

After the database is activated, subsequent device login is subject to the activated port bound WWN pairs.

Send documentation comments to mdsfeedback-doc@cisco.com.

When you activate the port security feature, the **auto-learn** option is also automatically enabled. You can choose to activate the port-security feature and disable **auto-learn** using the **port-security activate vsan number no-auto-learn** command.

To activate the port security feature, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# port-security activate vsan 1 | Activates the port security database for the specified VSAN, and automatically enables auto-learn. |
| | switch(config)# no port-security activate vsan 1 | Deactivates the port security database for the specified VSAN, and automatically disables auto-learn. |



Note

If required, you can disable autolearning (see the [“Disabling Autolearning”](#) section on page 21-7).

Database Activation Rejection

Database activation is rejected in the following cases:

- Missing or conflicting entries exist in the configuration database but not in the active database.
- If the auto-learn feature was enabled before the activation. To reactivate a database in this state.
- The exact security is not configured for each PortChannel member.
- The configured database is empty but the active database is not.

If the database activation is rejected due to one or more conflicts listed in the previous section, you may decide to proceed by forcing the port security activation.

Forceful Port Security Activation



Note

An activation using the **force** option can log out existing devices if they violate the active database.

You can view missing or conflicting entries using the **port-security database diff active vsan** command.

To forcefully activate the port security database, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# port-security activate vsan 1 force | Forces the VSAN 1 port security database to activate despite conflicts. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Database Reactivation



Tip If the **auto-learn** option is enabled and you activate the database, you will not be allowed to proceed.

To reactivate the database, follow these steps:

- Step 1** Disable auto-learning.
- Step 2** Copy the active database to the configured database.



Tip If the active database is empty, you cannot perform this step.

- Step 3** Make the required changes to the configuration database.
- Step 4** Activate the database.

To reactivate the port security database, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# no port-security auto-learn vsan 1 | Disables auto-learn and stops the switch from learning about new devices accessing the switch. Enforces the database contents based on the devices learned up to this point. |
| Step 3 | switch(config)# exit switch# port-security database copy vsan 1 | Copy from the active to the configured database. |
| Step 4 | switch# config t switch(config)# port-security activate vsan 1 | Activates the port security database for the specified VSAN, and automatically enables auto-learn. |

Send documentation comments to mdsfeedback-doc@cisco.com.

About AutoLearning

You can instruct the switch to automatically learn (auto-learn) the port security configurations over a specified period. This feature allows any switch in the Cisco MDS 9000 Family to automatically learn about devices and switches that connect to it. Use this feature to activate the port security feature for the first time as it saves tedious manual configuration for each port. You must configure the **auto-learn** option on a per-VSAN basis. If enabled, devices and switches that are allowed to connect to the switch are automatically learned, even if you have not configured any port access. Learned entries on a port are cleaned up after you shut down that port. Learning does not override the enforced port security policies.

When you activate the port security feature autolearning is also automatically enabled. When auto-learning is enabled, the following apply:

- Learning happens only for the devices or interfaces that were not activated.
- You will not be allowed to activate the database.

Configuring auto-learn

The state of the auto-learning configuration depends on the state of the port security feature:

- If the port security feature is not activated, the **auto-learn** option is disabled by default.
- If the port security feature is activated, the **auto-learn** option is enabled by default (unless you explicitly disabled this option).



Tip

If the **auto-learn** option is enabled on a VSAN, you can only activate the database for that VSAN by using the **force** option.

To enable auto-learning, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# port-security auto-learn vsan 1 | Enables auto-learn so the switch can learn about any device that is allowed to access VSAN 1. These devices are logged in the port security active database. |

Disabling Autolearning

To disable autolearning, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# no port-security auto-learn vsan 1 | Disables auto-learn and stops the switch from learning about new devices accessing the switch. Enforces the database contents based on the devices learned up to this point. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Auto-Learning Device Authorization

Table 21-1 summarizes the authorized connection for device requests.

Table 21-1 *Auto-Learn Device Authorization*

| Device (pWWN, nWWN, sWWN) | Requests Connection to | Authorization | Condition |
|--|--|---------------------------------|-----------|
| Configured with one or more switch ports | A switch on configured ports | Permitted | 1 |
| | A switch on other ports | Denied | 2 |
| Not configured | A port that is not configured | Permitted if auto-learn enabled | 3 |
| | | Denied if auto-learn disabled | 4 |
| Configured or not configured | A switch port that allows any device | Permitted | 5 |
| Configured to log in to any switch port | Any port on the switch | Permitted | 6 |
| Not configured | A port configured with some other device | Denied | 7 |

Authorization Scenario

Assume that the port security feature is activated and the following conditions are specified in the active database:

- A pWWN (P1) is allowed access through interface fc1/1 (F1).
- A pWWN (P2) is allowed access through interface fc1/1 (F1).
- A nWWN (N1) is allowed access through interface fc1/2 (F2).
- Any WWN is allowed access through interface fc1/3 (F3).
- A nWWN (N3) is allowed access through any interface.
- A pWWN (P3) is allowed access through interface fc1/4 (F4).
- A sWWN (S1) is allowed access through interface fc1/10-13 (F10 to F13).
- A pWWN (P10) is allowed access through interface fc1/11 (F11).

Table 21-2 summarizes the port security authorization results for this active database.

Table 21-2 *Authorization Results for Scenario*

| Scenario | Device Connection Request | Authorization | Condition | Reason |
|----------|---------------------------|---------------|-----------|------------------------|
| 1 | P1, N2, F1 | Permitted | 1 | No conflict. |
| 2 | P2, N2, F1 | Permitted | 1 | No conflict. |
| 3 | P3, N2, F1 | Denied | 2 | F1 is bound to P1/P2. |
| 4 | P1, N3, F1 | Permitted | 6 | Wildcard match for N3. |
| 5 | P1, N1, F3 | Permitted | 5 | Wildcard match for F3. |
| 6 | P1, N4, F5 | Denied | 2 | P1 is bound to F1. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 21-2 Authorization Results for Scenario (continued)

| Scenario | Device Connection Request | Authorization | Condition | Reason |
|----------|----------------------------|---------------|-----------|-------------------------------------|
| 7 | P5, N1, F5 | Denied | 2 | N1 is only allowed on F2. |
| 8 | P3, N3, F4 | Permitted | 1 | No conflict. |
| 9 | S1, F10 | Permitted | 1 | No conflict. |
| 10 | S2, F11 | Denied | 7 | P10 is bound to F11. |
| 11 | P4, N4, F5 (auto-learn on) | Permitted | 3 | No conflict. |
| 12 | P4, N4, F5(auto-learn off) | Denied | 4 | No match. |
| 13 | S3, F5 (auto-learn on) | Permitted | 3 | No conflict. |
| 14 | S3, F5 (auto-learn off) | Denied | 4 | No match. |
| 15 | P1, N1, F6 (auto-learn on) | Denied | 2 | P1 is bound to F1. |
| 16 | P5, N5, F1 (auto-learn on) | Denied | 7 | P3 is bound to F1. |
| 17 | S3, F4 (auto-learn on) | Denied | 7 | P3 paired with F4. |
| 18 | S1, F3 (auto-learn on) | Permitted | 5 | No conflict. |
| 19 | P5, N3, F3 | Permitted | 6 | Wildcard (*) match for F3 and N3. |
| 20 | P7, N3, F9 | Permitted | 6 | Wildcard (*) match for N3. |

Port Security Configuration Distribution

The port security feature uses the Cisco Fabric Services (CFS) infrastructure to enable efficient database management, provide a single point of configuration for the entire fabric in the VSAN, and enforce the port security policies on throughout the fabric (see [Chapter 9, “Using the CFS Infrastructure”](#)).

Enabling Distribution

All the configurations performed in distributed mode are stored in a pending (temporary) database. If you modify the configuration, you need to commit or discard the pending database changes to the configurations. The fabric remains locked during this period. Changes to the pending database are not reflected in the configurations until you commit the changes.

To enable the port security distribution, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# port-security distribute switch(config)# no port-security distribute | Enables distribution. Disables distribution. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Locking The Fabric

The first action that modifies the existing configuration creates the pending database and locks the feature in the VSAN. Once you lock the fabric, the following situations apply:

- No other user can make any configuration changes to this feature.
- A copy of the configuration database becomes the pending database.

Committing the Changes

If you commit the changes made to the configurations, the configurations in the pending database are distributed to other switches. On a successful commit, the configuration change is applied throughout the fabric and the lock is released.

To commit the port security configuration changes for the specified VSAN, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# port-security commit vsan 3 | Commits the port security changes in the specified VSAN. |

Discarding the Changes

If you discard (abort) the changes made to the pending database, the configurations remains unaffected and the lock is released.

To discard the port security configuration changes for the specified VSAN, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# port-security abort vsan 5 | Discards the port security changes in the specified VSAN and clears the pending configuration database. |

Activation and Autolearning Configuration Distribution

Activation and autolearning configurations in distributed mode are remembered merely as actions to be performed when you commit the changes in the pending database.

Learned entries are temporary and do not have any role in determining if a login is authorized or not. As such, learned entries do not participate in distribution. When you disable learning and commit the changes in the pending database, the learned entries become static entries in the active database and are distributed to all switches in the fabric. After the commit, the active database on all switches are identical and learning can be disabled.

If the pending database contains more than one activation and autolearning configuration when you commit the changes, then the activation and autolearning changes are consolidated and the behavior may change (see [Table 21-3](#)).

Send documentation comments to mdsfeedback-doc@cisco.com.

If you activate port security (the **port-security activate** command), follow up by disabling autolearning (the **no port-security auto-learn** command), and finally commit the changes in the pending database, then the net result of your actions is the same as issuing a **port-security activate no-auto-learn** command.

Table 21-3 **Scenarios for Activation and Autolearning Configurations in Distributed Mode**

| Scenario | Actions | Distribution = OFF | Distribution = ON |
|--|---|--|--|
| A and B exist in the configuration database, activation is not done and devices C,D are logged in. | 1. You activate the port security database and enable autolearning. | configuration database = {A,B} active database = {A,B, C ¹ , D*} | configuration database = {A,B} active database = {null} pending database = {A,B + activation to be enabled} |
| | 2. A new entry E is added to the configuration database. | configuration database = {A,B, E} active database = {A,B, C*, D*} | configuration database = {A,B} active database = {null} pending database = {A,B, E + activation to be enabled} |
| | 3. You issue a commit. | Not applicable | configuration database = {A,B, E} active database = {A,B, E, C*, D*} pending database = empty |
| A and B exist in the configuration database, activation is not done and devices C,D are logged in. | 1. You activate the port security database and enable autolearning. | configuration database = {A,B} active database = {A,B, C*, D*} | configuration database = {A,B} active database = {null} pending database = {A,B + activation to be enabled} |
| | 2. You disable learning. | configuration database = {A,B} active database = {A,B, C, D} | configuration database = {A,B} active database = {null} pending database = {A,B + activation to be enabled + learning to be disabled} |
| | 3. You issue a commit. | Not applicable | configuration database = {A,B} active database = {A,B} and devices C and D are logged out. This is equal to an activation with autolearning disabled. pending database = empty |

1. The * (asterisk) indicates learned entries.



Tip

In this case, we recommend that you perform a commit at the end of each operation: after you activate port security and after you enable auto learning.

Send documentation comments to mdsfeedback-doc@cisco.com.

Database Merge Guidelines

A database merge refers to a union of the configuration database and static (unlearned) entries in the active database. See the “[CFS Merge Support](#)” section on [page 9-7](#) for detailed concepts.

When merging the database between two fabric, follow these guidelines:

- Verify that the activation status and the auto-learn status is the same is both fabrics.
- Verify that the combined number of configuration for each VSAN in both databases does not exceed 2K.



If you do not follow these two conditions, the merge will fail. The next distribution will forcefully synchronize the databases and the activation states in the fabric.

Database Interaction

[Table 21-4](#) lists the differences and interaction between the active and configuration databases.

Table 21-4 Active and Configuration Port Security Databases

| Configuration Database | Active Database |
|---|---|
| Read-write. | Read-only. |
| Saving the configuration saves all the entries in the configuration database. | Saving the configuration only saves the activated entries. Learned entries are not saved. |
| Once activated, the configuration database can be modified without any effect on the active database. | Once activated, all devices that have already logged into the VSAN are also learned and added to the active database. |
| You can overwrite the configuration database with the active database. | You can overwrite the active database with the configured database by activating the port security database. Forcing an activation may violate the entries already configured in the active database. |



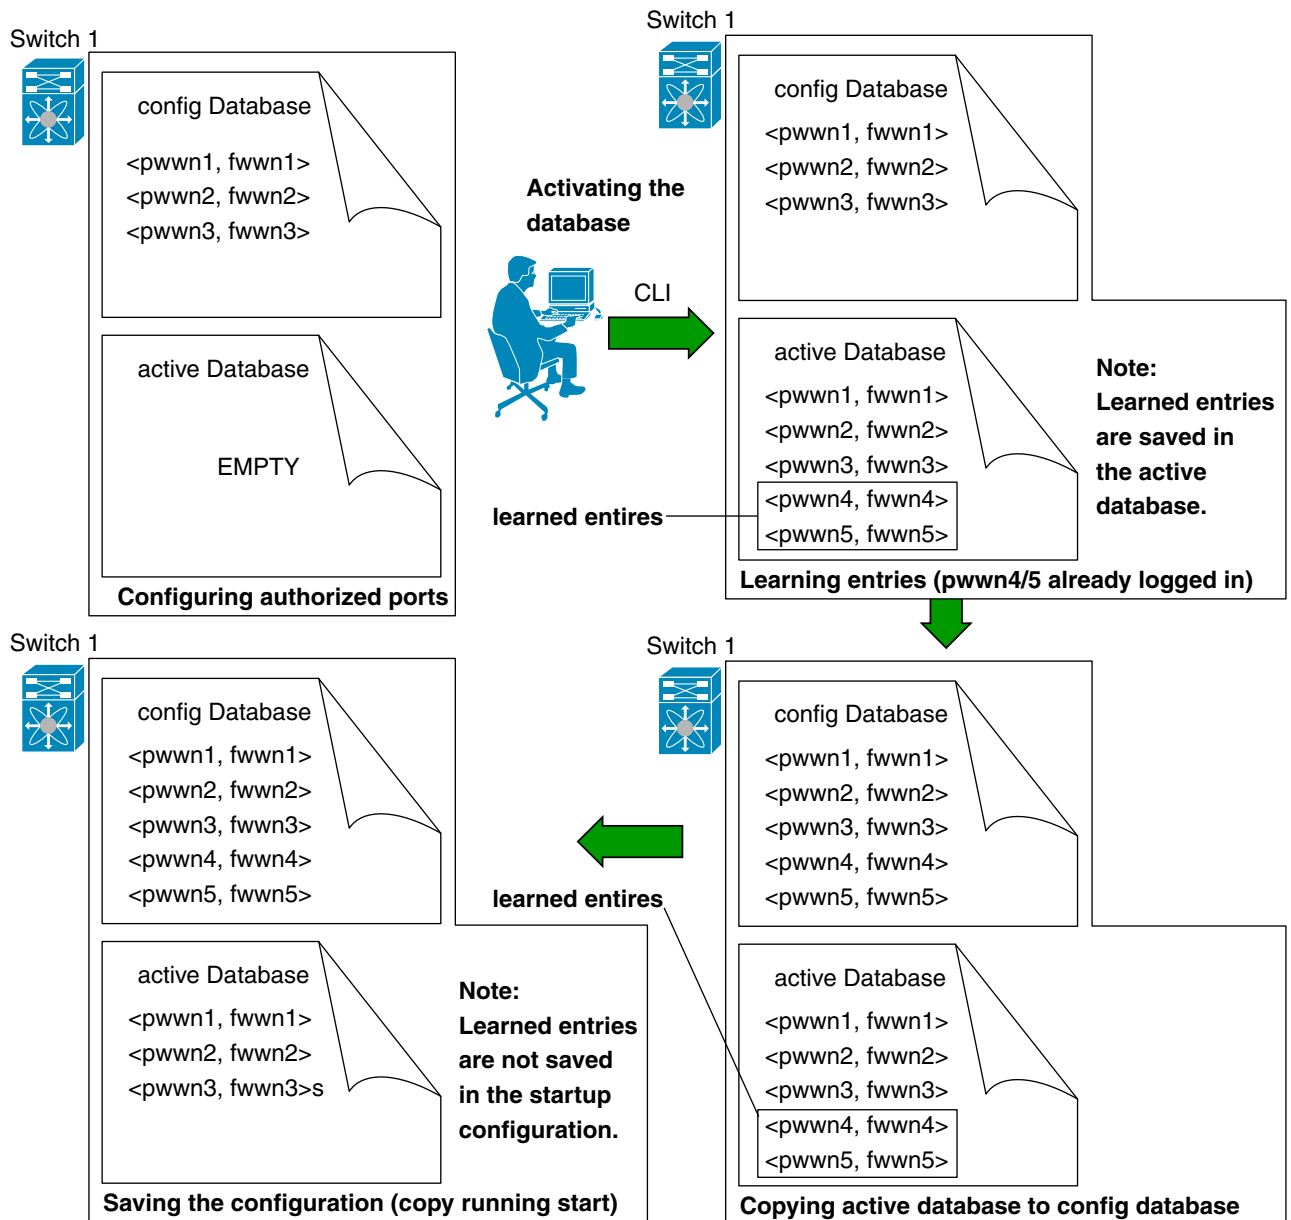
You can overwrite the configuration database with the active database using the **port-security database copy vsan** command. The **port-security database diff active vsan** command lists the differences between the active database and the configuration database.

Database Scenarios

[Figure 21-1](#) depicts various scenarios to depict the active database and the configuration database status based on port security configurations.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 21-1 Port Security Database Scenarios



Port Security Database Copy



Tip

We recommend that you issue **port-security database copy vsan** command after disabling autolearning. This action will ensure that the configuration database is in sync with the active database. If distribution is enabled, this command results in acquire of temporary copy (and consequently a fabric lock) of the configuration database. If you lock the fabric, you need to commit the changes to the configuration database of all the switches.

Send documentation comments to mdsfeedback-doc@cisco.com.

Use the **port-security database copy vsan** command to copy from the active to the configured database. If the active database is empty, this command is not accepted.

```
switch# port-security database copy vsan 1
switch#
```

Use the **port-security database diff active vsan** command to view the differences between the active database and the configuration database. This command can be used when resolving conflicts.

```
switch# port-security database diff active vsan 1
```

Use the **port-security database diff config vsan** command to obtain information on the differences between the configuration database and the active database.

```
switch# port-security database diff config vsan 1
```

Port Security Database Deletion



Tip

If the distribution is enabled, the deletion creates a copy of the database. An explicit **commit** command is required to actually delete the database.

Use the **no port-security** command in configuration mode to delete the configured database for a specified VSAN.

```
switch(config)# no port-security database vsan 1
```

Port Security Database Cleanup

Use the **clear port-security statistics** command to clear all existing statistics from the port security database for a specified VSAN.

```
switch# clear port-security statistics vsan 1
```

Use the **clear port-security database auto-learn interface** command to clear any learned entries in the active database for a specified interface within a VSAN.

```
switch# clear port-security database auto-learn interface fc1/1 vsan 1
```

Use the **clear port-security database auto-learn** command to clear any learned entries in the active database up to for the entire VSAN.

```
switch# clear port-security database auto-learn vsan 1
```



Note

The **clear port-security database auto-learn** and **clear port-security statistics** commands are only relevant to the local switch and do not acquire locks. Also, learned entries are only local to the switch and do not participate in distribution.

Use the **port-security clear vsan** command to clear the pending session in the VSAN from any switch in the VSAN.

```
switch# clear port-security session vsan 5
```


Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying Port Security Configurations

The **show port-security database** commands display the configured port security information (see Examples 21-1 to 21-11).

Example 21-1 *Displays the Contents of the Port Security Configuration Database*

```
switch# show port-security database
```

```
-----
VSAN      Logging-in Entity          Logging-in Point(      Interface)
-----
1         21:00:00:e0:8b:06:d9:1d(pwwn) 20:0d:00:05:30:00:95:de(fc1/13)
1         50:06:04:82:bc:01:c3:84(pwwn) 20:0c:00:05:30:00:95:de(fc1/12)
2         20:00:00:05:30:00:95:df(swgn) 20:0c:00:05:30:00:95:de(port-channel 128)
3         20:00:00:05:30:00:95:de(swgn) 20:01:00:05:30:00:95:de(fc1/1)
[Total 4 entries]
```

You can optionally specify a fWWN and a VSAN, or an interface and a VSAN in the **show port-security** command to view the output of the activated port security (see Example 21-2).

Example 21-2 *Displays the Port Security Configuration Database in VSAN 1*

```
switch# show port-security database vsan 1
```

```
-----
Vsan      Logging-in Entity          Logging-in Point      (Interface)
-----
1         *                    20:85:00:44:22:00:4a:9e (fc3/5)
1         20:11:00:33:11:00:2a:4a(pwwn) 20:81:00:44:22:00:4a:9e (fc3/1)
[Total 2 entries]
```

Example 21-3 *Displays the Activated Database*

```
switch# show port-security database active
```

```
-----
VSAN      Logging-in Entity          Logging-in Point(      Interface)      Learnt
-----
1         21:00:00:e0:8b:06:d9:1d(pwwn) 20:0d:00:05:30:00:95:de(fc1/13)      Yes
1         50:06:04:82:bc:01:c3:84(pwwn) 20:0c:00:05:30:00:95:de(fc1/12)      Yes
2         20:00:00:05:30:00:95:df(swgn) 20:0c:00:05:30:00:95:de(port-channel 128) Yes
3         20:00:00:05:30:00:95:de(swgn) 20:01:00:05:30:00:95:de(fc1/1)
[Total 4 entries]
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 21-4 *Displays the Contents of the Temporary Configuration Database*

```
switch# show port-security pending vsan 1
Session Context for VSAN 1
-----
Activation Status: Active
Auto Learn Status: On
Force activate: No
Config db modified: Yes
Activation done: Yes
Session owner: admin(2)
Session database:
-----
VSAN Logging-in Entity Logging-in Point (Interface)
-----
1 20:11:00:33:22:00:2a:4a(pwwn) 20:41:00:05:30:00:4a:1e(fc2/1)
[Total 1 entries]
```

Example 21-5 *Displays the Difference between the Temporary Configuration Database and the Configuration Database*

```
switch# show port-security pending-diff vsan 1
Session Diff for VSAN: 1
-----
Database will be activated
Learning will be turned ON
Database Diff:
+pwwn 20:11:00:33:22:00:2a:4a fwwn 20:41:00:05:30:00:4a:1e
```

The access information for each port can be individually displayed. If you specify the fwwn or interface options, all devices that are paired in the active database (at that point) with the given fWWN or the interface are displayed (see Examples 21-6 to 21-8).

Example 21-6 *Displays the Wildcard fWWN Port Security in VSAN 1*

```
switch# show port-security database fwwn 20:85:00:44:22:00:4a:9e vsan 1
Any port can login thru' this fwwn
```

Example 21-7 *Displays the Configured fWWN Port Security in VSAN 1*

```
switch# show port-security database fwwn 20:01:00:05:30:00:95:de vsan 1
20:00:00:0c:88:00:4a:e2(swwn)
```

Example 21-8 *Displays the Interface Port Information in VSAN 2*

```
switch# show port-security database interface fc 1/1 vsan 2
20:00:00:0c:88:00:4a:e2(swwn)
```

Send documentation comments to mdsfeedback-doc@cisco.com.

The port security statistics are constantly updated and available at any time (see [Example 21-9](#)).

Example 21-9 Displays the Port Security Statistics

```
switch# show port-security statistics
Statistics For VSAN: 1
-----
Number of pWWN permit: 2
Number of nWWN permit: 2
Number of sWWN permit: 2
Number of pWWN deny : 0
Number of nWWN deny : 0
Number of sWWN deny : 0

Total Logins permitted : 4
Total Logins denied   : 0
Statistics For VSAN: 2
-----
Number of pWWN permit: 0
Number of nWWN permit: 0
Number of sWWN permit: 2
Number of pWWN deny : 0
Number of nWWN deny : 0
Number of sWWN deny : 0
...
```

To verify the status of the active database and the auto-learn configuration, use the **show port-security status** command (see [Example 21-10](#)).

Example 21-10 Displays the Port Security Status

```
switch# show port-security status
Fabric Distribution Enabled
VSAN 1 :No Active database, learning is disabled, Session Lock Taken
VSAN 2 :No Active database, learning is disabled, Session Lock Taken
...
```

The **show port-security** command displays the previous 100 violations by default (see [Example 21-11](#)).

Example 21-11 Displays the Violations in the Port Security Database

```
switch# show port-security violations
```

| VSAN | Interface | Logging-in Entity | Last-Time | [Repeat count] |
|-------------------|----------------|-------------------------------|---------------------|----------------|
| 1 | fc1/13 | 21:00:00:e0:8b:06:d9:1d(pwwn) | Jul 9 08:32:20 2003 | [20] |
| | | 20:00:00:e0:8b:06:d9:1d(nwwn) | | |
| 1 | fc1/12 | 50:06:04:82:bc:01:c3:84(pwwn) | Jul 9 08:32:20 2003 | [1] |
| | | 50:06:04:82:bc:01:c3:84(nwwn) | | |
| 2 | port-channel 1 | 20:00:00:05:30:00:95:de(swwn) | Jul 9 08:32:40 2003 | [1] |
| [Total 2 entries] | | | | |

The **show port-security** command issued with the **last number** option displays only the specified number of entries that appear first.

Send documentation comments to mdsfeedback-doc@cisco.com.

Default Settings

Table 21-5 lists the default settings for all port security features in any switch.

Table 21-5 Default Security Settings

| Parameters | Default |
|--|--------------------------------------|
| Auto-learn | Enabled if port security is enabled. |
| Port security | Disabled. |
| Distribution | Disabled. |
| Note Enabling distribution enables it on all VSANs in the switch. | |



Configuring SNMP

From Cisco MDS SAN-OS Release 1.2, the CLI and SNMP use common roles in all switches in the Cisco MDS 9000 Family. You can use SNMP to modify a role that was created using CLI and vice versa.

From Cisco MDS SAN-OS Release 2.0(1b), users, passwords, and roles for all CLI and SNMP users are the same. A user configured through CLI can access the switch using the SNMP (for example, the Fabric Manager or the Device Manager) and vice versa.

This chapter includes the following sections:

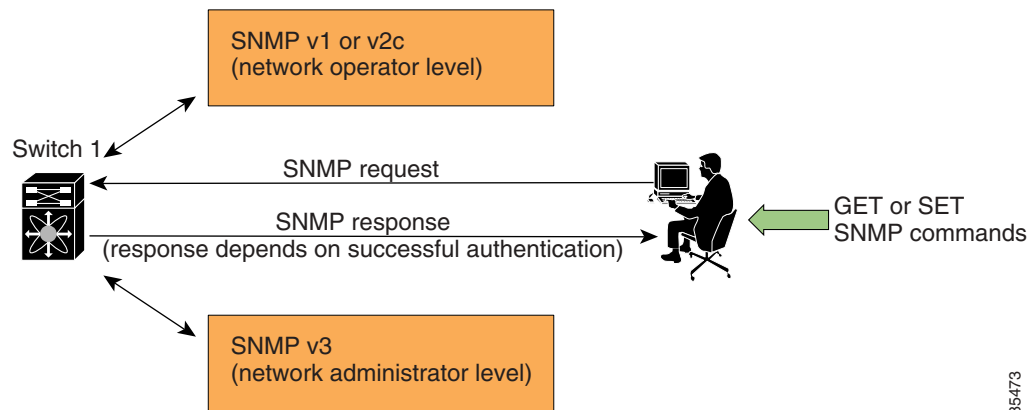
- [SNMP Security, page 22-2](#)
- [SNMPv3 CLI User Management and AAA Integration, page 22-3](#)
- [Restricting Switch Access, page 22-3](#)
- [Group-Based SNMP Access, page 22-4](#)
- [Configuring Common Roles, page 22-4](#)
- [Creating and Modifying Users, page 22-6](#)
- [Assigning SNMPv3 Users to Multiple Roles, page 22-8](#)
- [AES Encryption-Based Privacy, page 22-8](#)
- [Adding or Deleting Communities, page 22-9](#)
- [Assigning SNMP Switch Contact Information, page 22-9](#)
- [Configuring SNMP Notifications \(Traps and Informs\), page 22-9](#)
- [Displaying SNMP Security Information, page 22-13](#)
- [Default Settings, page 22-15](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

SNMP Security

SNMP is an application layer protocol that facilitates the exchange of management information between network devices. In all Cisco MDS 9000 Family switches, three SNMP versions are available: SNMPv1, SNMPv2c, and SNMPv3 (see Figure 22-1).

Figure 22-1 SNMP Security



85473

SNMP Version 1 and Version 2c

SNMP Version 1 (SNMPv1) and SNMP Version 2c (SNMPv2c) use a community string match for user authentication. Community strings provided a weak form of access control in earlier versions of SNMP. SNMPv3 provides much improved access control using strong authentication and should be preferred over SNMPv1 and SNMPv2c wherever it is supported.

SNMP Version 3

SNMP Version 3 (SNMPv3) is an interoperable standards-based protocol for network management. SNMPv3 provides secure access to devices by a combination of authenticating and encrypting frames over the network. The security features provided in SNMPv3 are:

- Message integrity—Ensures that a packet has not been tampered with in-transit.
- Authentication—Determines the message is from a valid source.
- Encryption—Scrambles the packet contents to prevent it from being seen by unauthorized sources.

SNMPv3 provides for both security models and security levels. A security model is an authentication strategy that is set up for a user and the role in which the user resides. A security level is the permitted level of security within a security model. A combination of a security model and a security level determines which security mechanism is employed when handling an SNMP packet.

Send documentation comments to mdsfeedback-doc@cisco.com.

SNMPv3 CLI User Management and AAA Integration

The Cisco SAN-OS software implement RFC 3414 and RFC 3415, including user-based security model (USM) and role-based access control. While SNMP and the CLI have common role management and share the same credentials and access privileges, the local user database was not synchronized in earlier releases.

As of Cisco SAN-OS Release 2.0(1b), SNMP v3 user management can be centralized at the AAA server level. This centralized user management allows the SNMP agent running on the Cisco MDS switch to leverage the user authentication service of AAA server. Once user authentication is verified, the SNMP PDUs are processed further. Additionally, the AAA server is also used to store user group names. SNMP uses the group names to apply the access/role policy that is locally available in the switch.

CLI and SNMP User Synchronization

Any configuration changes made to the user group, role, or password, results in the database synchronization for both SNMP and AAA.

To create an SNMP or CLI user, use either the **username** or **snmp-server user** commands.

- The `auth` passphrase specified in the **snmp-server user** command is synchronized as the password for the CLI user.
- The password specified in the **username** command is synchronized as the `auth` and `priv` passphrases for SNMP user.

Users are synchronized as follows:

- Deleting a user using either command results in the user being deleted for both SNMP and CLI.
- User-role mapping changes are synchronized in SNMP and CLI.



Note When the passphrase/password is specified in localized key/encrypted format, the password is not synchronized.

- Existing SNMP users continue to retain the `auth` and `priv` information without any changes.
- If a user is not present in one database and present in another database, the CLI user is created without any password (login is disabled) and the SNMP user is created with the `noAuthNoPriv` security level. Subsequently, the passwords and roles for these users will be synchronized.
- If the management station creates a SNMP user in the `usmUserTable`, the corresponding CLI user is created without any password (login is disabled) and will have the `network-operator` role.

Restricting Switch Access

You can restrict access to a Cisco MDS 9000 Family switch using IP Access Control Lists (IP-ACLs). See the [“IP Access Control Lists” section on page 26-5](#).

Send documentation comments to mdsfeedback-doc@cisco.com.

Group-Based SNMP Access



Note

Because *group* is a standard SNMP term used industry-wide, we refer to role(s) as group(s) in this SNMP section.

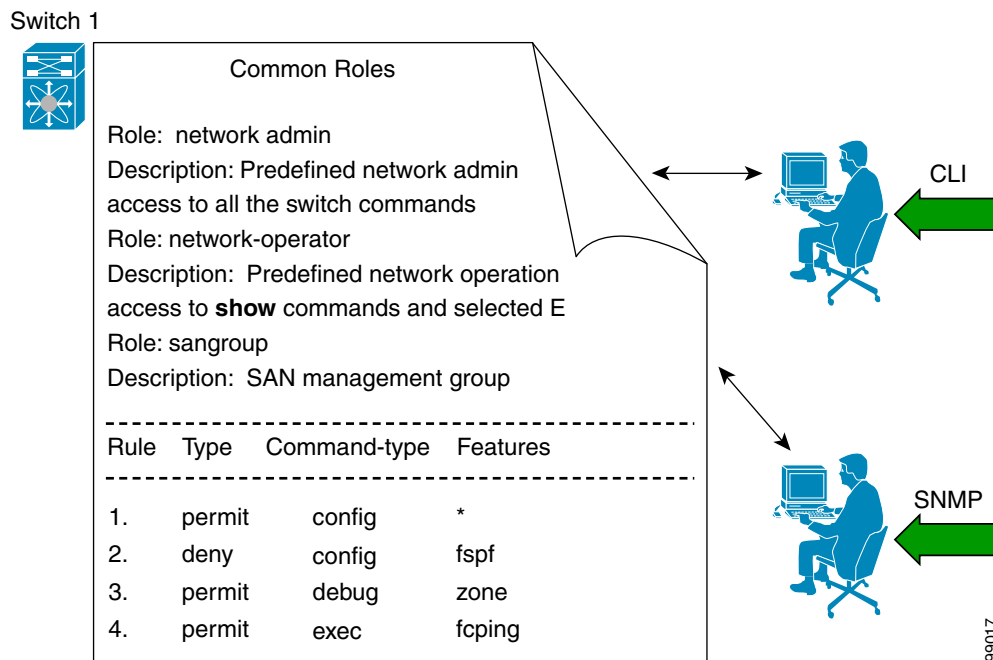
SNMP access rights are organized by groups. Each group in SNMP is similar to a role through the CLI. Each group is defined with three accesses: read access, write access, and notification access. Each access can be enabled or disabled within each group.

You can begin communicating with the agent once your user name is created, your roles are set up by your administrator, and you are added to the roles.

Configuring Common Roles

From Cisco MDS SAN-OS Release 1.2, CLI and SNMP in all switches in the Cisco MDS 9000 Family use common roles. You can use SNMP to modify a role that was created using CLI and vice versa (see [Figure 22-2](#)).

Figure 22-2 Common Roles



Each role in SNMP is the same as a role created or modified through the CLI (see the [“Role-Based Authorization”](#) section on page 19-21).

Each role can be restricted to one or more VSAN as required.

Send documentation comments to mdsfeedback-doc@cisco.com.

You can create new roles or modify existing roles using SNMP or the CLI.

- SNMP—Use the CISCO-COMMON-ROLES-MIB to configure or modify roles. Refer to the *Cisco MDS 9000 Family MIB Quick Reference*.
- CLI—Use the **role name** command.

Mapping of CLI operations to SNMP

SNMP has only three possible operations: GET, SET and NOTIFY. The CLI has five possible operations: DEBUG, SHOW, CONFIG, CLEAR and EXEC.



Note

NOTIFY does not have any restrictions like the syslog messages in the CLI.

Table 22-1 explains how the CLI operations are mapped to the SNMP operations.

Table 22-1 CLI Operation to SNMP Operation Mapping

| CLI Operation | SNMP Operation |
|---------------|----------------|
| DEBUG | Ignored |
| SHOW | GET |
| CONFIG | SET |
| CLEAR | SET |
| EXEC | SET |

Example

The following example shows the privileges and rules mapping CLI operations to SNMP operations for a role named my_role.

```
switch# show role name my_role
Role:my_role
  vsan policy:permit (default)
-----
Rule      Type      Command-type      Feature
-----
1.  permit    clear              *
2.  deny      clear              ntp
3.  permit    config             *
4.  deny      config             ntp
5.  permit    debug              *
6.  deny      debug              ntp
7.  permit    show               *
8.  deny      show               ntp
9.  permit    exec               *
```



Note

Although CONFIG is denied for NTP in rule 4, rule 9 allows the SET to NTP MIB objects because EXEC also maps to the SNMP SET operation.

Send documentation comments to mdsfeedback-doc@cisco.com.

Creating and Modifying Users

You can create users or modify existing users using SNMP or the CLI.

- **SNMP**—Create a user as a clone of an existing user in the `usmUserTable` on the switch. Once you have created the user, change the cloned secret key before activating the user. Refer to RFC 2574.
- **CLI**—Create a user or modify an existing user using the **`snmp-server user`** command.

By default only two roles are available in a Cisco MDS 9000 Family switch—`network-operator` and `network-admin`. You can also use any role that is configured in the Common Roles database (see the “[Configuring Common Roles](#)” section on page 22-4).



Tip

As of Cisco MDS SAN-OS Release 2.0(1b), all updates to the CLI security database and the SNMP user database are synchronized. After upgrading to Release 2.0(1b), you can use the SNMP password to log into either Fabric Manager or Device Manager. However, after you use the CLI password to log into Fabric Manager or Device Manager, you must use the CLI password for all future logins. If a user exists in both the SNMP database and the CLI database before upgrading to Release 2.0(1b), then the set of roles assigned to the user becomes the union of both sets of roles after the upgrade.

Configuring SNMP Users from the CLI

As of Release 2.0(1b), the passphrase specified in **`snmp-server user`** command and the **`username`** command are synchronized (see the “[SNMPv3 CLI User Management and AAA Integration](#)” section on page 22-3).

To create or modify SNMP users from the CLI, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | <code>switch# config t</code> | Enters configuration mode. |
| Step 2 | <code>switch(config)# snmp-server user joe network-admin auth sha abcd1234</code> | Creates or modifies the settings for a user (joe) in the <code>network-admin</code> role using the HMAC-SHA-96 authentication password (abcd1234). |
| | <code>switch(config)# snmp-server user sam network-admin auth md5 abcdefgh</code> | Creates or modifies the settings for a user (sam) in the <code>network-admin</code> role using the HMAC-MD5-96 authentication password (abcdefgh). |
| | <code>switch(config)# snmp-server user Bill network-admin auth sha abcd1234 priv abcdefgh</code> | Creates or modifies the settings for a user (network-admin) in the <code>network-admin</code> role using the HMAC-SHA-96 authentication level and privacy encryption parameters. |
| | <code>switch(config)# no snmp-server user usernameA</code> | Deletes the user (usernameA) and all associated parameters. |
| | <code>switch(config)# no snmp-server usam role vsan-admin</code> | Deletes the specified user (usam) from the <code>vsan-admin</code> role. |
| | <code>switch(config)# snmp-server user user1 network-admin auth md5 0xab0211gh priv 0x45abf342 localizedkey</code> | Specifies the password to be in localized key format (see RFC 2574). The localized key is provided in the hex format (for example, 0xacbdef). |
| | <code>switch(config)# snmp-server user user2 auth md5 asdgfsadf priv aes-128 asgfsgkhkj</code> | Configures the user2 with the MD5 authentication protocol and AES-128 privacy protocol. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|--|---|
| Step 3 | switch(config)# snmp-server user joe sangroup | Adds the specified user (joe) to the sangroup role. |
| | switch(config)# snmp-server user joe techdocs | Adds the specified user (joe) to the techdocs role. |



Caution

Avoid using the **localizedkey** option when configuring an SNMP user from the CLI. The localized keys are not portable across devices as they contain device engine ID information. If a configuration file is copied to the device, the passwords may not be set correctly if the configuration file was generated at a different device. Explicitly configure the desired passwords after copying the configuration into the device. Passwords specified with the **localizedkey** option are limited to 130 characters.



Note

The **snmp-server user** command takes engineID as an additional parameter. The engineID is for creating the notification target user (see to the [“Configuring the Notification Target User”](#) section on page 22-12). If the engineID is not specified, the local user is created.

Enforcing SNMPv3 Message Encryption

By default the SNMP agent allows the securityLevel parameters of 'authNoPriv' and 'authPriv' for the SNMPv3 messages that use SNMPv3 user configured with 'auth' and 'priv' keys. You can enforce the message encryption for a user by using the following configuration commands:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# snmp-server user testUser enforcePriv | Enforces the message encryption for SNMPv3 messages using this user. |
| | switch(config)# no snmp-server user testUser enforcePriv | Disables SNMPv3 message encryption enforcement. |

Note You can only use this command for previously existing users configured with both auth and priv keys. When the user is configured to enforce privacy, for any SNMPv3 PDU request using such a user with securityLevel parameter of either 'noAuthNoPriv' or 'authNoPriv', the SNMP agent responds with 'authorizationError'.

Alternatively, you can enforce the SNMPv3 message encryption globally on all the users using the following commands:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# snmp-server globalEnforcePriv | Enforces the SNMPv3 message encryption for all the users on the switch. |
| | switch(config)# no snmp-server globalEnforcePriv | Disables global SNMPv3 message encryption enforcement. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Assigning SNMPv3 Users to Multiple Roles

As of Cisco SAN-OS Release 2.0(1b), the SNMP server user configuration is enhanced to accommodate multiple roles (groups) for SNMPv3 users. After the initial SNMPv3 user creation, you can map additional roles for the user.



Note

Only users belonging to network-admin role can assign roles to other users.

To configure multiple roles for SNMPv3 users from the CLI, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# snmp-server user NewUser role1 | Creates or modifies the settings for an SNMPv3 user (NewUser) for the role1 role. |
| | switch(config)# snmp-server user NewUser role2 | Creates or modifies the settings for an SNMPv3 user (NewUser) for the role2 role. |
| | switch(config)# no snmp-server user User5 role2 | Removes role2 for the specified user (User5) |

AES Encryption-Based Privacy

The Advanced Encryption Standard (AES) is the symmetric cipher algorithm. The Cisco SAN-OS software uses AES as one of the privacy protocols for SNMP message encryption and conforms with RFC3826.

As of Cisco SAN-OS Release 2.0(1b), the **priv** option offers a choice of DES or 128-bit AES encryption for SNMP security encryption. The **priv** option along with **aes-128** token indicates that this privacy password is for generating 128-bit AES key. The AES priv password can have a minimum of eight characters. If the passphrases are specified in clear text, you can specify a maximum of 64 characters. If you use the localized key, you can specify a maximum of 130 characters.



Note

For an SNMPv3 operation using the external AAA server, user configurations in the external AAA server require AES to be the privacy protocol to use SNMP PDU encryption.

To create or modify SNMP users from the CLI, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# snmp-server user user1 role1 auth md5 0xab0211gh priv des 0x45abf342 localizedkey | Specifies the password to be in localized key format using the DES option for security encryption |
| | switch(config)# snmp-server user user1 role2 auth sha 0xab0211gh priv aes-128 0x45abf342 localizedkey | Specifies the password to be in localized key format using the 128-bit AES option for security encryption |

Send documentation comments to mdsfeedback-doc@cisco.com.

Adding or Deleting Communities

You can configure read-only or read-write access for SNMPv1 and SNMPv2 users. Refer to RFC 2576.

To create an SNMPv1 or SNMPv2c community, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# snmp-server community snmp_Community ro | Adds read-only access for the specified SNMP community. |
| | switch(config)# snmp-server community snmp_Community rw | Adds read-write access for the specified SNMP community. |
| | switch(config)# no snmp-server community snmp_Community | Deletes access for the specified SNMP community (default). |

Assigning SNMP Switch Contact Information

The switch contact information is limited to 32 characters (without spaces).

Use the **snmp-server** command to set the contact information and the switch location. Use the **no** form of the command to remove the system contact information.

To configure contact information, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# snmp-server contact NewUser | Assigns the contact name for the switch. |
| | switch(config)# no snmp-server contact NewUser | Deletes the contact name for the switch. |
| Step 3 | switch(config)# snmp-server location SanJose | Assigns the switch location. |
| | switch(config)# no snmp-server location SanJose | Deletes the switch location. |

Configuring SNMP Notifications (Traps and Informs)

You can configure the Cisco MDS switch using the CLI to send notifications to SNMP managers when particular events occur. You can send these notifications as Traps or Informs. Traps are unreliable because the receiver does not send any acknowledgment when it receives a trap. The sender cannot determine if the trap was received. However, an SNMP manager that receives Informs acknowledges the message with an SNMP Response PDU. If the sender never receives a Response, the inform is normally retransmitted. Thus, Informs are more likely to reach their intended destination.



Note

Use the SNMP-TARGET-MIB to obtain more information on the destinations to which notifications are to be sent either as Traps or as Informs. Refer to the *Cisco MDS 9000 Family MIB Quick Reference* for more information.



Tip

The SNMP version 1 option is not available with the **snmp-server host ip-address informs** command.

Send documentation comments to mdsfeedback-doc@cisco.com.

Configuring SNMPv1 and SNMPv2c Notifications

To configure SNMPv1 and SNMPv2c notifications, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# snmp-server host 171.71.187.101 traps version 2c private udp-port 1163 | Configures the specified host to receive SNMPV2c trap notifications. |
| | switch(config)# no snmp-server host 171.71.187.101 traps version 2c private udp-port 2162 | Prevents the specified host from receiving SNMPv2c trap notifications on the configured UDP port. |
| Step 3 | switch(config)# snmp-server host 171.71.187.101 informs version 2c private udp-port 1163 | Configures the specified host to receive SNMPV2c inform notifications using SNMPv2c community string “private”. |
| | switch(config)# no snmp-server host 171.71.187.101 informs version 2c private udp-port 2162 | Prevents the specified host from receiving SNMP v2c inform notifications on the configured UDP port. |

Configuring SNMPv3 Notifications

To configure SNMPv3 notifications, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# snmp-server host 16.20.11.14 traps version 3 noauth testuser udp-port 1163 | Configures the specified host to receive SNMPv3 trap notifications using SNMPv3 user “testuser” and securityLevel of “noAuthNoPriv”. |
| | switch(config)# snmp-server host 16.20.11.14 informs version 3 auth testuser udp-port 1163 | Configures the specified host to receive SNMPv3 inform notifications using SNMPv3 user “testuser” and securityLevel of “AuthNoPriv”. |
| | switch(config)# snmp-server host 16.20.11.14 informs version 3 priv testuser udp-port 1163 | Configures the specified host to receive SNMPv3 inform notifications using SNMPv3 user “testuser” and securityLevel of “AuthPriv”. |
| | switch(config)# no snmp-server host 172.18.2.247 informs version 3 testuser noauth udp-port 2162 | Prevents the specified host from receiving SNMPv3 inform notifications. |



Note

In the case of SNMPv3 trap notifications, the SNMP manager is expected to know the user credentials (authKey/PrivKey) based on the switch’s engineID to authenticate and decrypt the SNMP messages.

Send documentation comments to mdsfeedback-doc@cisco.com.

Enabling SNMP Notifications

Notifications (Traps and Informs) are system alerts that the switch generates when certain events occur. By default, no notification is defined or issued. If a notification name, is not specified all notifications are disabled or enabled.

Table 22-2 lists the trap notifications that are disabled by default. This list does not include the **entity fru**, **vrrp**, **license**, unlisted trap notifications, and other generic trap notifications such as **coldstart**, **warmstart**, **linkup**, and **linkdown**.

Table 22-2 List of SNMP Trap Notifications Enabled by Default

| Traps enabled | Related Commands |
|--------------------------------|---|
| All traps listed in this table | snmp-server enable traps |
| Entity traps | snmp-server enable traps entity snmp-server enable traps entity fru |
| FCC trap | snmp-server enable traps fcc |
| FC domain trap | snmp-server enable traps fcdomain |
| FC name server trap | snmp-server enable traps fcns |
| FCS trap | snmp-server enable traps fcs discovery-complete snmp-server enable traps fcs request-reject |
| FDMI trap | snmp-server enable traps fdmi |
| FSPF trap | snmp-server enable traps fspf |
| License manager trap | snmp-server enable traps license |
| Port security trap | snmp-server enable traps port-security |
| RSCN traps | snmp-server enable traps rscn snmp-server enable traps rscn els snmp-server enable traps rscn ils |
| SNMP agent traps | snmp-server enable traps snmp snmp-server enable traps snmp authentication |
| VRRP trap | snmp-server enable traps vrrp |
| Zone traps | snmp-server enable traps zone snmp-server enable traps zone default-zone-behavior-change snmp-server enable traps zone merge-failure snmp-server enable traps zone merge-success snmp-server enable traps zone request-reject snmp-server enable traps zone unsupp-mem |

To enable trap notification, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# snmp-server enable traps fcdomain | Enables the specified SNMP trap (fcdomain) notification. |
| | switch(config)# no snmp-server enable traps | Disables the specified SNMP trap notification. If a trap name is not specified, all traps are disabled. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Use the **show snmp trap** command to display all the traps and their status.

```
switch# show snmp trap
Trap type                               Enabled
-----
entity fru                             Yes
fcc                                    No
fcdomain                              No
fcns                                   No
fcs request-reject                     No
fcs discovery-complete                 No
fdmi                                   No
fspf                                   No
license                               Yes
port-security                          No
rscn els                              No
rscn ils                              No
snmp authentication                   No
vrrp                                  Yes
zone unsupported member                No
zone request-reject                   No
zone merge-failure                    No
zone merge-success                    No
zone default-zone-behavior-change     No
```

Configuring the Notification Target User

You must configure a notification target user on the switch for sending SNMPv3 inform notifications to the SNMP manager.

To configure the notification target user, use the following command:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# snmp-server user testusr auth md5 xyub20gh priv xyub20gh engineID 00:00:00:63:00:01:00:a1:ac:15:10:03 | Configures the notification target user with the specified credentials for the SNMP manager with the specified engineID |
| | switch(config)# no snmp-server user testusr auth md5 xyub20gh priv xyub20gh engineID 00:00:00:63:00:01:00:a1:ac:15:10:03 | Removes the notification target user. |

The credentials of the notification target user are used for encrypting the SNMPv3 inform notification messages to the configured SNMP manager (as in the **snmp-server host** command)



Note

For authenticating and decrypting the received INFORM PDU, the SNMP manager should have the same user credentials in its local configuration data store of users.

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying SNMP Security Information

Use the **show snmp** commands to display configured SNMP information (see [Example 22-1](#) and [22-3](#)).

Example 22-1 Displays SNMP User Details

```
switch# show snmp user
```

| SNMP USERS | | | |
|------------|------|---------------|--------------------|
| User | Auth | Priv(enforce) | Groups |
| admin | md5 | des(no) | network-admin |
| testusr | md5 | aes-128(no) | role111 role222 |

```
NOTIFICATION TARGET USERS (configured for sending V3 Inform)
```

| User | Auth | Priv |
|--|------|------|
| testtargetusr (EngineID 0:0:0:63:0:1:0:0:0:15:10:3) | md5 | des |

Example 22-2 Displays SNMP Community Information

```
switch# show snmp community
```

| Community | Access |
|------------|--------|
| ----- | ----- |
| private | rw |
| public | ro |
| v93RACqPNH | ro |

Example 22-3 Displays SNMP Host Information

```
switch# show snmp host
```

| Host | Port | Version | Level | Type | SecName |
|---------------|------|---------|--------|------|---------|
| 171.16.126.34 | 2162 | v2c | noauth | trap | public |
| 171.16.75.106 | 2162 | v2c | noauth | trap | public |
| ... | | | | | |
| 171.31.58.97 | 2162 | v2c | auth | trap | public |
| ... | | | | | |

Send documentation comments to mdsfeedback-doc@cisco.com.

The **show snmp** command displays counter information for SNMP contact, location, and packet settings. This command provides information that is used entirely by the Cisco MDS 9000 Family Fabric Manager (refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*). See [Example 22-4](#).

Example 22-4 Displays SNMP Information

```
switch# show snmp
sys contact:
sys location:
1631 SNMP packets input
    0 Bad SNMP versions
    0 Unknown community name
    0 Illegal operation for community name supplied
    0 Encoding errors
    64294 Number of requested variables
    1 Number of altered variables
    1628 Get-request PDUs
    0 Get-next PDUs
    1 Set-request PDUs
152725 SNMP packets output
    0 Too big errors
    1 No such name errors
    0 Bad values errors
    0 General errors

Community                               Group / Access
-----
public                                   rw
```

| SNMP USERS | | | |
|------------|------|---------------|--------------------|
| User | Auth | Priv(enforce) | Groups |
| admin | md5 | des(no) | network-admin |
| testusr | md5 | aes-128(no) | role111 role222 |

NOTIFICATION TARGET USERS (configured for sending V3 Inform)

| User | Auth | Priv |
|---------------|------|------|
| testtargetusr | md5 | des |

Send documentation comments to mdsfeedback-doc@cisco.com.

```
(EngineID 0:0:0:63:0:1:0:0:0:15:10:3)
```

Example 22-5 Displays SNMP Engine IDs

```
switch# show snmp engineID
Local SNMP engineID: 800000090300053000851E
```

Example 22-6 Displays Information on SNMP Security Groups

```
switch# show snmp group
groupname: network-admin
security model: any
security level: noAuthNoPriv
readview: network-admin-rd
writeview: network-admin-wr
notifyview: network-admin-rd
storage-type: permanent
row status: active

groupname: network-admin
security model: any
security level: authNoPriv
readview: network-admin-rd
writeview: network-admin-wr
notifyview: network-admin-rd
storage-type: permanent
row status: active

groupname: network-operator
security model: any
security level: noAuthNoPriv
readview: network-operator-rd
writeview: network-operator-wr
notifyview: network-operator-rd
storage-type: permanent
row status: active

groupname: network-operator
security model: any
security level: authNoPriv
readview: network-operator-rd
writeview: network-operator-wr
notifyview: network-operator-rd
storage-type: permanent
row status: active
```

Default Settings

Table 22-3 lists the default settings for all SNMP features in any switch.

Table 22-3 Default SNMP Settings

| Parameters | Default |
|--------------|--------------------------------|
| User account | No expiry (unless configured). |
| Password | None. |

Send documentation comments to mdsfeedback-doc@cisco.com.



CHAPTER 23

Configuring RMON

RMON is an Internet Engineering Task Force (IETF) standard monitoring specification that allows various network agents and console systems to exchange network monitoring data. You can use the RMON alarms and events to monitor Cisco MDS 9000 Family switches running the Cisco SAN-OS Release 2.0(1b) or higher software.

This chapter includes the following sections:

- [About RMON, page 23-1](#)
- [Configuring RMON, page 23-1](#)
- [RMON Verification, page 23-3](#)
- [Default Settings, page 23-3](#)

About RMON

All switches in the Cisco MDS 9000 Family support the following RMON functions (defined in RFC 2819):

- **Alarm**—Monitors a specific management information base (MIB) object for a specified interval, triggers an alarm at a specified value (rising threshold), and resets the alarm at another value (falling threshold). Alarms can be used with events; the alarm triggers an event, which can generate a log entry or an SNMP trap.
- **Event**—Determines the action to take when an event is triggered by an alarm. The action can be to generate a log entry, an SNMP trap, or both.

Refer to the *Cisco MDS 9000 Family MIB Quick Reference* for agent and management information.

Refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide* for information on an SNMP-compatible network management station.

See the “[SNMP Security](#)” section on [page 19-32](#) for SNMP security-related CLI configurations.

Configuring RMON

RMON is disabled by default and no events or alarms are configured in the switch. You can configure your RMON alarms and events by using the CLI or an SNMP-compatible network management station.

Send documentation comments to mdsfeedback-doc@cisco.com.


Tip

We recommend an additional, generic RMON console application on the network management station (NMS) to take advantage of RMON's network management capabilities. Refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*.


Note

You must also configure SNMP on the switch to access RMON MIB objects.

RMON Alarm Configuration

You can set an alarm on any MIB object. The specified MIB must be an existing SNMP MIB object in standard dot notation (1.3.6.1.2.1.2.2.1.14.16777216 for ifInOctets.16777216).

Use one of the following options to specify the interval to monitor the MIB variable (ranges from 1 to 4294967295 seconds):

- Use the **delta** option to test the change between samples of a MIB variable.
- Use the **absolute** option to test each MIB variable directly.
- Use the **delta** option to test any MIB objects that are counters.

The range for the **rising threshold** and **falling threshold** values is -2147483647 to 2147483647.


Caution

The **falling threshold** must be less than the **rising threshold**.

You can optionally specify the following parameters:

- The event-number to trigger if the rising or falling threshold exceeds the specified limit.
- The owner of the alarm.

To enable RMON alarms, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# rmon alarm 20 1.3.6.1.2.1.2.2.1.14.16777216 2900 delta rising-threshold 15 1 falling-threshold 0 owner test | Configures RMON alarm number 20 to monitor the 1.3.6.1.2.1.2.2.1.14.16777216 once every 900 seconds until the alarm is disabled and checks the change in the variable's rise or fall. If the value shows a MIB counter increase of 15 or more, the software triggers an alarm. The alarm in turn triggers event number 1, which is configured with the RMON event command. Possible events can include a log entry or an SNMP trap. If the MIB value changes by 0, the alarm is reset and can be triggered again. |
| | switch(config)# no rmon alarm 2 | Deletes the specified entry from the alarm table |

Send documentation comments to mdsfeedback-doc@cisco.com.

RMON Event Configuration

To enable RMON events, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# rmon event 2 log trap eventtrap description CriticalErrors owner Test2 | Creates RMON event number 2 to define CriticalErrors and generates a log entry when the event is triggered by the alarm. The user Test2 owns the row that is created in the event table by this command. This example also generates an SNMP trap when the event is triggered. |
| | switch(config)# no rmon event 5 | Deletes an entry from the RMON event table. |

RMON Verification

Use the **show rmon** and **show snmp** commands to display configured RMON and SNMP information (see [Example 23-1](#) and [23-2](#)).

Example 23-1 Displays Configured RMON Alarms

```
switch# show rmon alarms
Alarm 1 is active, owned by admin
Monitors 1.3.6.1.2.1.2.2.1.16.16777216 every 1 second(s)
Taking delta samples, last value was 0
Rising threshold is 1, assigned to event 0
Falling threshold is 0, assigned to event 0
On startup enable rising or falling alarm
```

Example 23-2 Displays Configured RMON Events

```
switch# show rmon events
Event 2 is active, owned by Test2
Description is CriticalErrors
Event firing causes log and trap to community eventtrap, last fired 0
Event 500 is active, owned by admin
Description is
Event firing causes log, last fired 138807208
```

Default Settings

[Table 23-1](#) lists the default settings for all RMON features in any switch.

Table 23-1 Default RMON Settings

| Parameters | Default |
|-------------|-----------|
| RMON alarms | Disabled. |
| RMON events | Disabled. |

Send documentation comments to mdsfeedback-doc@cisco.com.



Configuring Fibre Channel Routing Services and Protocols

Fabric Shortest Path First (FSPF) is the standard path selection protocol used by Fibre Channel fabrics. The FSPF feature is enabled by default on all Fibre Channel switches. Except in configurations that require special consideration, you do not need to configure any FSPF services. FSPF automatically calculates the best path between any two switches in a fabric. Specifically, FSPF is used to:

- Dynamically compute routes throughout a fabric by establishing the shortest and quickest path between any two switches.
- Select an alternative path in the event of the failure of a given path. FSPF supports multiple paths and automatically computes an alternative path around a failed link. It provides a preferred route when two equal paths are available.

This chapter provides details on Fibre Channel routing services and protocols. It includes the following sections:

- [FSPF Features, page 24-2](#)
- [FSPF Examples, page 24-2](#)
- [FSPF Global Configuration, page 24-4](#)
- [FSPF Interface Configuration, page 24-6](#)
- [Configuring Fibre Channel Routes, page 24-8](#)
- [Clearing FSPF Counters, page 24-9](#)
- [Broadcast and Multicast Routing, page 24-10](#)
- [Broadcast and Multicast Routing, page 24-10](#)
- [In-Order Delivery, page 24-11](#)
- [Flow Statistics Configuration, page 24-15](#)
- [Displaying Routing and Forwarding Information, page 24-17](#)
- [Default Settings, page 24-21](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

FSPF Features

FSPF is the protocol currently standardized by the T11 committee for routing in Fibre Channel networks. The FSPF protocol has the following characteristics and features:

- Supports multipath routing.
- Bases path status on a link state protocol.
- Routes hop by hop, based only on the domain ID.
- Runs only on E ports or TE ports and provides a loop free topology.
- Runs on a per VSAN basis. Connectivity in a given VSAN in a fabric is guaranteed only for the switches configured in that VSAN.
- Uses a topology database to keep track of the state of the links on all switches in the fabric and associates a cost with each link.
- Guarantees a fast reconvergence time in case of a topology change. Uses the standard Dijkstra's algorithm, but there is a static dynamic option for a more robust, efficient, and incremental Dijkstra's algorithm. The reconvergence time is fast and efficient as the route computation is done on a per VSAN basis.

FSPF Examples

This section provides examples of topologies and applications that demonstrate the benefits of FSPF.




Note

The FSPF feature can be used on any topology.

Fault Tolerant Fabric

Figure 24-1 depicts a fault tolerant fabric using a partial mesh topology. If a link goes down anywhere in the fabric, any switch can still communicate with all others in the fabric. In the same way, if any switch goes down, the connectivity of the rest of the fabric is preserved.

Figure 24-1 Fault Tolerant Fabric



For example, if all links are of equal speed, the FSPF calculates two equal paths from A to C: A-D-C (green) and A-E-C (blue).


Send documentation comments to mdsfeedback-doc@cisco.com.

Redundant Links

To further improve on the topology in Figure 24-1, each connection between any pair of switches can be replicated; two or more links can be present between a pair of switches. Figure 24-2 shows this arrangement. Because switches in the Cisco MDS 9000 Family support PortChanneling, each pair of physical links can appear to the FSPF protocol as one single logical link.

By bundling pairs of physical links, FSPF efficiency is considerably improved by the reduced database size and the frequency of link updates. Once physical links are aggregated, failures are not attached to a single link but to the entire PortChannel. This configuration also improves the resiliency of the network. The failure of a link in a PortChannel does not trigger a route change, thereby reducing the risks of routing loops, traffic loss, or fabric downtime for route reconfiguration.

Figure 24-2 Fault Tolerant Fabric with Redundant Links




For example, if all links are of equal speed and no PortChannels exist, the FSPF calculates four equal paths from A to C: A1-E-C, A2-E-C, A3-D-C, and A4-D-C. If PortChannels exist, these paths are reduced to two.

Fail-Over Scenarios for PortChannels and FSPF Links

The SmartBits traffic generator was used to evaluate the scenarios displayed in Figure 24-3. Two links between switch 1 and switch 2 exist as either equal-cost ISLs or PortChannels. There is one flow from traffic generator 1 to traffic generator 2. The traffic was tested at 100% utilization of 1 Gbps in two scenarios:

- Disabling the traffic link by either physically removing the cable (see Table 24-1).
- Shutting down either switch 1 or switch 2 (see Table 24-2).

Figure 24-3 Fail-Over Scenario Using Traffic Generators



Send documentation comments to mdsfeedback-doc@cisco.com.

Table 24-1 *Physically Removing the Cable for the SmartBits Scenario*

| PortChannel Scenario | | FSPF Scenario (Equal cost ISL) | |
|---|----------|--------------------------------|----------|
| Switch 1 | Switch 2 | Switch 1 | Switch 2 |
| 110 ms (~2K frame drops) | | 130 ms+ (~4k frame drops) | |
| 100 ms (hold time when a signal loss is reported as mandated by the standard) | | | |

Table 24-2 *Shutting Down the Switch for the SmartBits Scenario*

| PortChannel Scenario | | FSPF Scenario (Equal cost ISL) | |
|------------------------|--------------------------|--------------------------------|-------------------------|
| Switch 1 | Switch 2 | Switch 1 | Switch 2 |
| ~0 ms (~8 frame drops) | 110 ms (~2K frame drops) | 130 ms+ (~4K frame drops) | |
| No hold time needed | Signal loss on switch 1 | No hold time needed | Signal loss on switch 1 |

FSPF Global Configuration

By default, FSPF is enabled on switches in the Cisco MDS 9000 Family.

Some FSPF features can be globally configured in each VSAN. By configuring a feature for the entire VSAN, you do not have to specify the VSAN number for every command. This global configuration feature also reduces the chance of typing errors or other minor configuration errors.



Note

FSPF is enabled by default. Generally, you do not need to configure these advanced features.



Caution

The default for the backbone region is 0 (zero). You do not need to change this setting unless your region is different from the default. If you are operating with other vendors using the backbone region, you can change this default to be compatible with those settings.

Global FSPF Configuration

To configure a FSPF feature for the entire VSAN, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fspf config vsan 1 | Enters FSPF global configuration mode for the specified VSAN. |
| Step 3 | switch-config-(fspf-config)# spf static | Forces static SPF computation for the dynamic (default) incremental VSAN. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|---|---|
| Step 4 | switch-config-(fspf-config) # spf hold-time 10 | Configures the hold time between two route computations in milliseconds (ms) for the entire VSAN. The default value is 0. Note If the specified time is shorter, the routing is faster. However, the processor consumption increases accordingly. |
| Step 5 | switch-config-(fspf-config) # region 7 | Configures the autonomous region for this VSAN and specifies the region ID (7). |

FSPF Configuration Deletion

To delete the FSPF configuration for the entire VSAN, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# no fspf config vsan 3 | Deletes the FSPF configuration for VSAN 3. |

FSPF Routing Protocol Usage

To enable or disable FSPF routing protocols, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# no fspf enable vsan 5 | Disables the FSPF routing protocol in VSAN 5. |
| | switch(config)# fspf enable vsan 7 | Enables the FSPF routing protocol in VSAN 7. |

Link State Record Defaults

Each time a new switch enters the fabric, a link state record (LSR) is sent to the neighboring switches, and then flooded throughout the fabric. [Table 24-3](#) displays the default settings for switch responses.

Table 24-3 LSR Default Settings

| LSR Option | Default | Description |
|--|------------|---|
| Acknowledgment interval (RxmtInterval) | 5 seconds | The time a switch waits for an acknowledgment from the LSR before retransmission. |
| Refresh time (LSRefreshTime) | 30 minutes | The time a switch waits before sending an LSR refresh transmission. |
| Maximum age (MaxAge) | 60 minutes | The time a switch waits before dropping the LSR from the database. |

Send documentation comments to mdsfeedback-doc@cisco.com.

FSPF Interface Configuration

Several FSPF commands are available on a per interface basis. These configuration procedures apply to an interface in a specific VSAN.

FSPF Link Cost

FSPF tracks the state of links on all switches in the fabric, associates a cost with each link in its database, and then chooses the path with a minimal cost. The cost associated with an interface can be administratively changed to implement the FSPF route selection. The integer value to specify cost can range from 1 to 65,535. The default cost for 1 Gbps is 1000 and for 2Gbps is 500.

To configure FSPF link cost, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/4 switch(config-if)# | Configures the specified interface, or if already configured, enters configuration mode for the specified interface. |
| Step 3 | switch(config-if)# fspf cost 5 vsan 90 | Configures the cost for the selected interface in VSAN 90. |

Hello Time Intervals

You can set the FSPF Hello time interval to specify the interval between the periodic hello messages sent to verify the health of the link. The integer value can range from 1 to 65,535 seconds.



Note

This value must be the same in the ports at both ends of the ISL.

To configure the FSPF Hello time interval, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/4 switch(config-if)# | Configures the specified interface, or if already configured, enters configuration mode for the specified interface. |
| Step 3 | switch(config-if)# fspf hello-interval 15 vsan 175 switch(config-if)# | Specifies the hello message interval (15 seconds) to verify the health of the link in VSAN 175. The default is 20 seconds. |

Dead Time Intervals

You can set the FSPF dead time interval to specify the maximum interval for which a hello message must be received before the neighbor is considered lost and removed from the database. The integer value can range from 1 to 65,535 seconds.

Send documentation comments to mdsfeedback-doc@cisco.com.

**Note**

This value must be the same in the ports at both ends of the ISL.

**Caution**

An error is reported at the command prompt if the configured dead time interval is less than the hello time interval.

To configure the FSPF dead time interval, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/4 switch(config-if)# | Configures the specified interface, or if already configured, enters configuration mode for the specified interface. |
| Step 3 | switch(config-if)# fspf dead-interval 25 vsan 7 switch(config-if)# | Specifies the maximum interval for VSAN 7 before which a hello message must be received on the selected interface before the neighbor is considered lost. The default is 80 seconds. |

Disabling FSPF for Specific Interfaces

You can disable the FSPF protocol for selected interfaces. By default, FSPF is enabled on all E ports and TE ports. This default can be disabled by setting the interface as passive.

**Note**

FSPF must be enabled at both ends of the interface for the protocol to work.

To disable FSPF for a specific interface, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/4 switch(config-if)# | Configures a specified interface, or if already configured, enters configuration mode for the specified interface. |
| Step 3 | switch(config-if)# fspf passive vsan 1 switch(config-if)# | Disables the FSPF protocol for the specified interface in the specified VSAN. |
| | switch(config-if)# no fspf passive vsan 1 switch(config-if)# | Reenables the FSPF protocol for the specified interface in the specified VSAN. |

Retransmitting Intervals

You can specify the time after which an unacknowledged link state update should be transmitted on the interface. The integer value to specify retransmit intervals can range from 1 to 65,535 seconds.

**Note**

This value must be the same on the switches on both ends of the interface.

Send documentation comments to mdsfeedback-doc@cisco.com.


To configure the FSPF retransmit time interval, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/4 switch(config-if)# | Configures the specified interface, or if already configured, enters configuration mode for the specified interface. |
| Step 3 | switch(config-if)# fspf retransmit-interval 15 vsan 12 switch(config-if)# | Specifies the retransmit time interval for unacknowledged link state updates in VSAN 12. The default is 5 seconds. |

Configuring Fibre Channel Routes

Each port implements forwarding logic, which forwards frames based on its FC ID. To configure the FC ID for the specified interface and domain, you can configure the specified route (for example FC ID 111211 and domain ID 3) in the switch with domain ID 1 (see [Figure 24-4](#)).

Figure 24-4 Fibre Channel Routes



Note

Other than in VSANs, run-time checks are not performed on configured and suspended static routes.

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure an FC route, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# fcroute 0x111211 interface fc1/1 domain 3 vsan 2 switch(config)# | Configures the route for the specified Fibre Channel interface and domain. In this example, interface fc1/1 is assigned an FC ID (0x111211) and a domain ID (3) to the next hop switch. |
| | switch(config)# fcroute 0x111211 interface port-channel 1 domain 3 vsan 4 switch(config)# | Configures the route for the specified PortChannel interface and domain. In this example, interface port-channel 1 is assigned an FC ID (0x111211) and a domain ID (3) to the next hop switch. |
| | switch(config)# fcroute 0x031211 interface fc1/1 domain 3 metric 1 vsan 1 switch(config-if)# | Configures the static route for a specific FC ID and next hop domain ID and also assigns the cost of the route. If the remote destination option is not specified, the default is direct. |
| | switch(config)# fcroute 0x111112 interface fc1/1 domain 3 metric 3 remote vsan 3 | Adds a static route to the RIB. If this is an active route and the FIB ¹ records are free, it is also added to the FIB. If the cost (metric) of the route is not specified, the default is 10. |
| Step 3 | switch(config)# fcroute 0x610000 0xff0000 interface fc 1/1 domain 1 vsan 2 switch(config)# | Configures the netmask for the specified route in interface fc1/1 (or PortChannel). You can specify one of three routes: ff0000 matches only the domain, ffff00 matches the domain and the area, ffffff matches the domain, area, and port. |

1. FIB = Forwarding Information Base

Clearing FSPF Counters

To clear the FSPF statistics counters for one interface or for the entire VSAN, follow this step:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# clear fspf counters vsan 1 switch# | Clears the FSPF statistics counters for the specified VSAN. If an interface reference is not specified, all counters are cleared. |
| | switch# clear fspf counters vsan 200 interface fc1/1 switch# | Clears the FSPF statistics counters for the specified interface in VSAN 200. |

[Send documentation comments to mdsfeedback-doc@cisco.com.](mailto:mdsfeedback-doc@cisco.com)

Broadcast and Multicast Routing

Broadcast and Multicast in a Fibre Channel fabric uses the concept of a distribution tree to reach all switches in the fabric.

FSPF provides the topology information to compute the distribution tree. Fibre Channel defines 256 multicast groups and one broadcast address for each VSAN. Switches in the Cisco MDS 9000 Family only use broadcast routing. By default, they use the principal switch as the root node to derive a loop-free distribution tree for multicast and broadcast routing in a VSAN.



Caution

All switches in the fabric should run the same multicast and broadcast distribution tree algorithm to ensure the same distribution tree.

Prior to Cisco SAN-OS Release 2.0(1b), the SAN-OS software used the principal switch to compute the multicast tree.

As of Cisco SAN-OS Release 2.0(1b) to interoperate with other vendor switches (following FC-SW3 guidelines), the SAN-OS software uses the lowest domain switch as the root to compute the multicast tree in interop mode.

By default, the **native** (non-interop) mode uses the principal switch as the root. If you change the default, be sure to configure the same mode in all switches in the fabric. Otherwise, multicast traffic could face potential loop and frame-drop problems.



Note

The operational mode can be different from the configured interop mode. The interop mode always uses the lowest domain switch as the root.

Use the **mcast root lowest vsan** command to change the multicast root from the principal switch to lowest domain switch.

To use the lowest domain switch for the multicast tree computation, follow these steps.

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# mcast root lowest vsan 1 | Uses the lowest domain switch to compute the multicast tree. |
| | switch(config)# mcast root principal vsan 1 | Defaults to using the principal switch to compute the multicast tree. |

To display the configured and operational multicast mode and the selected root domain, use the **show mcast** command.

```
switch# show mcast vsan 1
Multicast root for VSAN 1
    Configured root mode : Principal switch
    Operational root mode : Principal switch
    Root Domain ID : 0xef(239)
```

Send documentation comments to mdsfeedback-doc@cisco.com.

In-Order Delivery

In-Order Delivery (IOD) of data frames guarantees frame delivery to a destination in the same order that they were sent by the originator.

Some Fibre Channel protocols or applications cannot handle out-of-order frame delivery. In these cases, switches in the Cisco MDS 9000 Family preserve frame ordering in the frame flow. The source ID (SID), destination ID (DID), and optionally the originator exchange ID (OX ID) identify the flow of the frame.

In case of a single switch, all frames received by a specific ingress port and destined to a certain egress port are always delivered in the same order in which they were received.




Tip

If you enable the IOD feature, the graceful shutdown feature is not implemented.

Reordering Network Frames

When you experience a route change in the network, the new selected path may be faster or less congested than the old route (see [Figure 24-5](#)).

Figure 24-5 Route Change Delivery



In [Figure 24-5](#), the new path from Switch 1 to Switch 4 is faster. Hence, Frame 3 and Frame 4 may be delivered before Frame 1 and Frame 2.

If the in-order guarantee feature is enabled, the frames within the network are treated as follows:


- Frames in the network are delivered in the order in which they are transmitted.
- Frames that cannot be delivered in order within the network latency drop period are dropped inside the network.

Reordering PortChannel Frames

When a link change occurs in a PortChannel, the frames for the same exchange or the same flow can switch from one path to another faster path (see [Figure 24-6](#)).

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 24-6 Link Congestion Delivery



In Figure 24-6, the port of the old path (red dot) is congested. Hence Frame 3 and Frame 4 can be delivered before Frame 1 and Frame 2.

When the in-order guarantee feature is enabled, the frames crossing a PortChannel are treated as follows:

- Frames using the old path are delivered before new frames are accepted.
- Frames that cannot be delivered in order through the old path within the switch latency drop period are dropped.
- The new frames are delivered through the new path after the switch latency drop period has elapsed.

Enabling In-Order Delivery

By default, in-order delivery is disabled on switches in the Cisco MDS 9000 Family.

As of Cisco MDS SAN-OS Release 1.3(4), you can enable the in-order delivery feature for a specific VSAN or for the entire switch.



Tip

We recommend that you only enable this feature when devices that cannot handle any out-of-order frames are present in the switch. Load-balancing algorithms within the Cisco MDS 9000 Family ensure that frames are delivered in order during normal fabric operation. The load-balancing algorithms based on source FC ID, destination FC ID, and exchange ID are enforced in hardware without any performance degradation. However, if the fabric encounters a failure and this feature is enabled, the recovery will be delayed because of an intentional pausing of fabric forwarding to purge the fabric of resident frames that could potentially be forwarded out-of-order.

Enabling IOD Globally

To verify that the IOD parameters are uniform across all VSANs, enable IOD globally before performing a downgrade to a Cisco MDS SAN-OS Release 1.3(3) or earlier.

Issue the **in-order-guarantee** command to ensure that the IOD parameters are uniform across VSANs.

To enable in-order delivery for the switch, follow these steps.

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# in-order-guarantee | Enables in-order delivery in the switch. |
| | switch(config)# no in-order-guarantee | Reverts the switch to the factory defaults and disables the in-order delivery feature. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Enabling IOD for a VSAN

When you create a new VSAN, that VSAN automatically inherits the global in-order-guarantee value.

You can subsequently change the in-order-guarantee for the new VSAN using the **in-order-guarantee vsan** command. You can display both the global and per VSAN values using the **show in-order-guarantee** command

To enable in-order delivery for a specific VSAN, follow these steps.

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# in-order-guarantee vsan 3452 switch(config)# no in-order-guarantee 101 | Enables in-order delivery in VSAN 3452. Reverts the switch to the factory defaults and disables the in-order delivery feature. |

Displaying the IOD Status

Use the **show in-order-guarantee** command to display the present configuration status:

```
switch# show in-order-guarantee
global inorder delivery configuration:guaranteed

VSAN specific settings
vsan 1 inorder delivery:guaranteed
vsan 101 inorder delivery:not guaranteed
vsan 1000 inorder delivery:guaranteed
vsan 1001 inorder delivery:guaranteed
vsan 1682 inorder delivery:guaranteed
vsan 2001 inorder delivery:guaranteed
vsan 2009 inorder delivery:guaranteed
vsan 2456 inorder delivery:guaranteed
vsan 3277 inorder delivery:guaranteed
vsan 3451 inorder delivery:guaranteed
vsan 3452 inorder delivery:guaranteed
vsan 3453 inorder delivery:guaranteed
```

[Send documentation comments to mdsfeedback-doc@cisco.com.](mailto:mdsfeedback-doc@cisco.com)

Configuring the Drop Latency Time

Use the **fcdroplateny network** command to change the default latency time for either a network, a specified VSAN in a network, or for the entire switch.

To configure the network and the switch drop latency time, follow these steps.

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fcdroplateny network 5000 | Configures network drop latency time to be 5000 ms for the network. The valid range is 0 to 60000 ms. The default is 2000 ms. Note The network drop latency must be computed as the sum of all switch latencies of the longest path in the network |
| | switch(config)# fcdroplateny network 6000 vsan 3 | Configures network drop latency time to be 6000 ms for VSAN 3. |
| | switch(config)# no fcdroplateny network 4500 | Removes the current fcdroplateny network configuration (4500) and reverts the switch to the factory defaults. |
| Step 3 | switch(config)# fcdroplateny switch 4000 | Configures switch drop latency time to be 4000 ms for the switch. The valid range is 0 to 60000 ms. The default is 500 ms. Note The switch drop latency parameter should have the same value in all the switches in the network |
| | switch(config)# no fcdroplateny switch 4500 | Removes the current fcdroplateny switch configuration (4500) and reverts the switch to the factory defaults. |

Displaying Latency Information

You can view the configured latency parameters using the **show fcdroplateny** command (see [Example 24-1](#)).

Example 24-1 Displays Administrative Distance

```
switch# show fcdroplateny
switch latency value:500 milliseconds
global network latency value:2000 milliseconds
```

```
VSAN specific network latency settings
vsan 1 network latency:5000 milliseconds
vsan 2 network latency:2000 milliseconds
vsan 103 network latency:2000 milliseconds
vsan 460 network latency:500 milliseconds
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Flow Statistics Configuration

Flow statistics count the ingress traffic in the aggregated statistics table. You can collect two kinds of statistics:

- Aggregated flow statistics to count the traffic for a VSAN.
- Flow statistics to count the traffic for a source and destination ID pair in a VSAN.

If you enable flow counters, you can enable a maximum of 1K entries for aggregate flow and flow statistics. Be sure to assign an unused flow index to a module for each new flow. Flow indexes can be repeated across modules. The number space for flow index is shared between the aggregate flow statistics and the flow statistics.

Configuring Flow Statistics

To count the aggregated flow statistics for a VSAN, follow these steps:

| | Command | Purpose |
|--------|---|---------------------------------------|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fcflow stats aggregated module 1 index 1005 vsan 1 switch(config)# | Enables the aggregated flow counter. |
| | switch(config)# no fcflow stats aggregated module 1 index 1005 vsan 1 switch(config)# | Disables the aggregated flow counter. |

Counting Flow Statistics

To count the flow statistics for a source and destination FC ID in a VSAN, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fcflow stats module 1 index 1 0x145601 0x5601ff ffffffff vsan 1 switch(config)# | Enables the flow counter. Note The source ID and the destination ID are specified in FC ID hex format (for example, 0x123aff). The mask can be one of ff0000 or ffffffff. |
| Step 3 | switch(config)# no fcflow stats aggregated module 2 index 1001 vsan 2 switch(config)# | Disables the flow counter. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Clearing FIB Statistics

Use the **clear fcflow stats** command to clear the aggregated flow counter (see Examples 24-2 and 24-3).

Example 24-2 Clears Aggregated Flow Counters

```
switch# clear fcflow stats aggregated module 2 index 1
```

Example 24-3 Clears Flow Counters for Source and Destination FC IDs

```
switch# clear fcflow stats module 2 index 1
```

Displaying Flow Statistics

Use the **show fcflow stats** commands to view flow statistics (see Example 24-4 to 24-6).

Example 24-4 Displays Aggregated fcflow Details for the Specified Module

```
switch# show fcflow stats aggregated module 2
Idx  VSAN # frames # bytes
----  -
0000 4    387,653 674,235,875
0001 6    34,402  2,896,628
```

Example 24-5 Displays fcflow Details for the Specified Module

```
switch# show fcflow stats module 2
Idx  VSAN D ID      S ID      mask      # frames # bytes
----  -
0000 4    032.001.002 007.081.012 ff.ff.ff   387,653 674,235,875
0001 6    004.002.001 019.002.004 ff.00.00   34,402  2,896,628
```

Example 24-6 Displays fcflow Index Usage for the Specified Module

```
switch# show fcflow stats usage module 2
2 flows configured
configured flow : 3,7
```


Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying Routing and Forwarding Information

You can view specific information about existing Fibre Channel and FSPF configurations (see Examples 24-7 to 24-15).

Example 24-7 Displays Administrative Distance

```
switch# show fcroute distance
```

| UUID | Route Distance | Name |
|------|----------------|------------|
| 10 | 20 | RIB |
| 22 | 40 | FCDOMAIN |
| 39 | 80 | RIB-CONFIG |
| 12 | 100 | FSPF |
| 17 | 120 | FLOGI |
| 21 | 140 | TLP |
| 14 | 180 | MCAST |
| 64 | 200 | RIB-TEST |



Note

When the number of routes are displayed in the command output, both visible and hidden routes are included in the total number of routes. While hidden routes are added to the count, they are not visible.

Example 24-8 Displays Multicast Routing Information

```
switch# show fcroute multicast
```

| VSAN | FC ID | # Interfaces |
|------|------------|--------------|
| 1 | 0xffffffff | 0 |
| 2 | 0xffffffff | 1 |
| 3 | 0xffffffff | 1 |
| 4 | 0xffffffff | 0 |
| 5 | 0xffffffff | 0 |
| 6 | 0xffffffff | 0 |
| 7 | 0xffffffff | 0 |
| 8 | 0xffffffff | 0 |
| 9 | 0xffffffff | 0 |
| 10 | 0xffffffff | 0 |

Example 24-9 Displays FCID Information for a Specified VSAN

```
switch# show fcroute multicast vsan 3
```

| VSAN | FC ID | # Interfaces |
|------|------------|--------------|
| 3 | 0xffffffff | 1 |

Example 24-10 Displays FCID and interface Information for a Specified VSAN

```
switch# show fcroute multicast 0xffffffff vsan 2
```

| VSAN | FC ID | # Interfaces |
|------|------------|--------------|
| 2 | 0xffffffff | 1 |
| | fc1/1 | |

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 24-11 Displays Unicast Routing Information

```
switch# show fcroute unicast
D:direct R:remote P:permanent V:volatile A:active N:non-active
# Next
Protocol VSAN    FC ID/Mask      RCtrl/Mask  Flags Hops  Cost
-----
static  1      0x010101 0xffffffff 0x00 0x00 D P A 1      10
static  2      0x111211 0xffffffff 0x00 0x00 R P A 1      10
fspf    3      0x610000 0xff0000 0x00 0x00 D P A 4     500
static  4      0x040101 0xffffffff 0x00 0x00 R P A 1     103
static  4      0x040102 0xffffffff 0x00 0x00 R P A 1     103
static  4      0x040103 0xffffffff 0x00 0x00 R P A 1     103
static  4      0x040104 0xffffffff 0x00 0x00 R P A 1     103
static  4      0x111211 0xffffffff 0x00 0x00 D P A 1     10
```

Example 24-12 Displays Unicast Routing Information for a Specified VSAN

```
switch# show fcroute unicast vsan 4
D:direct R:remote P:permanent V:volatile A:active N:non-active
# Next
Protocol VSAN    FC ID/Mask      RCtrl/Mask  Flags Hops  Cost
-----
static  4      0x040101 0xffffffff 0x00 0x00 R P A 1     103
static  4      0x040102 0xffffffff 0x00 0x00 R P A 1     103
static  4      0x040103 0xffffffff 0x00 0x00 R P A 1     103
static  4      0x040104 0xffffffff 0x00 0x00 R P A 1     103
static  4      0x111211 0xffffffff 0x00 0x00 D P A 1     10
```

Example 24-13 Displays Unicast Routing Information for a Specified FCID

```
switch# show fcroute unicast 0x040101 0xffffffff vsan 4
D:direct R:remote P:permanent V:volatile A:active N:non-active
# Next
Protocol VSAN    FC ID/Mask      RCtrl/Mask  Flags Hops  Cost
-----
static  4      0x040101 0xffffffff 0x00 0x00 R P A 1     103
      fcl/2 Domain 0xa6(166)
```

Example 24-14 Displays Route Database Information

```
switch# show fcroute summary
FC Route Database Created Thu Feb 13 07:21:52 2003
VSAN      Ucast      Mcast      Label      Last Modified Time
-----
1          5          1          0          Thu Feb 13 10:21:06 2003
2          4          1          0          Thu Feb 13 10:21:07 2003
3          4          1          0          Thu Feb 13 10:21:08 2003
4          4          1          0          Thu Feb 13 10:21:09 2003
5          4          1          0          Thu Feb 13 10:21:10 2003
6          4          1          0          Thu Feb 13 10:21:11 2003
7          4          1          0          Thu Feb 13 10:21:12 2003
8          4          1          0          Thu Feb 13 10:21:13 2003
9          4          1          0          Thu Feb 13 10:21:14 2003
10         4          1          0          Thu Feb 13 10:21:15 2003
11         4          1          0          Thu Feb 13 10:21:16 2003
12         4          1          0          Thu Feb 13 10:21:17 2003
13         4          1          0          Thu Feb 13 10:21:18 2003
14         4          1          0          Thu Feb 13 10:21:18 2003
15         4          1          0          Thu Feb 13 10:21:19 2003
-----
Total      61          15          0
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 24-15 Displays Route Database Information for a Specified VSAN

```
switch# show fcroute summary vsan 5
FC Route Database Created Thu Feb 13 07:21:52 2003
```

| VSAN | Ucast | Mcast | Label | Last Modified Time |
|-------|-------|-------|-------|--------------------------|
| 5 | 4 | 1 | 0 | Thu Feb 13 10:21:10 2003 |
| Total | 4 | 1 | 0 | |

Displaying Global FSPF Information

Example 24-16 displays global FSPF information for a specific VSAN:

- Domain number of the switch.
- Autonomous region for the switch.
- Min_LS_arrival: minimum time that must elapse before the switch accepts LSR updates.
- Min_LS_interval: minimum time that must elapse before the switch can transmit an LSR.



Tip If the Min_LS_interval is higher than 10 seconds, the graceful shutdown feature is not implemented.

- LS_refresh_time: interval time lapse between refresh LSR transmissions.
- Max_age: maximum time aa LSR can stay before being deleted.

Example 24-16 Displays FSPF Information for a Specified VSAN

```
switch# show fspf vsan 1
FSPF routing for VSAN 1
FSPF routing administration status is enabled
FSPF routing operational status is UP
It is an intra-domain router
Autonomous region is 0
SPF hold time is 0 msec
MinLsArrival = 1000 msec , MinLsInterval = 5000 msec
Local Domain is 0x65(101)
Number of LSRs = 3, Total Checksum = 0x0001288b

Protocol constants :
  LS_REFRESH_TIME = 1800 sec
  MAX_AGE         = 3600 sec

Statistics counters :
  Number of LSR that reached MaxAge = 0
  Number of SPF computations         = 7
  Number of Checksum Errors          = 0
  Number of Transmitted packets :   LSU 65 LSA 55 Hello 474 Retranmsitted LSU 0
  Number of received packets :     LSU 55 LSA 60 Hello 464 Error packets 10
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying the FSPF Database

Example 24-17 displays a summary of the FSPF database for a specified VSAN. If other parameters are not specified, all LSRs in the database are displayed:

- LSR type
- Domain ID of the LSR owner
- Domain ID of the advertising router
- LSR age
- LSR incarnation member
- Number of links

You could narrow the display to obtain specific information by issuing additional parameters for the domain ID of the LSR owner. For each interface, the following information is also available:

- Domain ID of the neighboring switch
- E port index
- Port index of the neighboring switch
- Link type and cost

Example 24-17 Displays FSPF Database Information

```
switch# show fspf database vsan 1
```

```
FSPF Link State Database for VSAN 1 Domain 0x0c(12)
```

```
LSR Type           = 1
Advertising domain ID = 0x0c(12)
LSR Age            = 1686
LSR Incarnation number = 0x80000024
LSR Checksum       = 0x3caf
Number of links     = 2
```

| NbrDomainId | IfIndex | NbrIfIndex | Link Type | Cost |
|-------------|------------|------------|-----------|------|
| 0x65(101) | 0x0000100e | 0x00001081 | 1 | 500 |
| 0x65(101) | 0x0000100f | 0x00001080 | 1 | 500 |

```
FSPF Link State Database for VSAN 1 Domain 0x65(101)
```

```
LSR Type           = 1
Advertising domain ID = 0x65(101)
LSR Age            = 1685
LSR Incarnation number = 0x80000028
LSR Checksum       = 0x8443
Number of links     = 6
```

| NbrDomainId | IfIndex | NbrIfIndex | Link Type | Cost |
|-------------|------------|------------|-----------|------|
| 0xc3(195) | 0x00001085 | 0x00001095 | 1 | 500 |
| 0xc3(195) | 0x00001086 | 0x00001096 | 1 | 500 |
| 0xc3(195) | 0x00001087 | 0x00001097 | 1 | 500 |
| 0xc3(195) | 0x00001084 | 0x00001094 | 1 | 500 |
| 0x0c(12) | 0x00001081 | 0x0000100e | 1 | 500 |
| 0x0c(12) | 0x00001080 | 0x0000100f | 1 | 500 |

Send documentation comments to mdsfeedback-doc@cisco.com.

```
FSPF Link State Database for VSAN 1 Domain 0xc3(195)
LSR Type = 1
Advertising domain ID = 0xc3(195)
LSR Age = 1686
LSR Incarnation number = 0x80000033
LSR Checksum = 0x6799
Number of links = 4
  NbrDomainId      IfIndex  NbrIfIndex  Link Type  Cost
-----
  0x65(101) 0x00001095    0x00001085        1      500
  0x65(101) 0x00001096    0x00001086        1      500
  0x65(101) 0x00001097    0x00001087        1      500
  0x65(101) 0x00001094    0x00001084        1      500
```

Displaying FSPF Interfaces

[Example 24-18](#) displays the following information for each selected interface.

- Link cost
- Timer values
- Neighbor's domain ID (if known)
- Local interface number
- Remote interface number (if known)
- FSPF state of the interface
- Interface counters

Example 24-18 Displays FSPF Interface Information

```
switch# show fspf vsan 1 interface fc1/1
FSPF interface fc1/1 in VSAN 1
FSPF routing administrative state is active
Interface cost is 500
Timer intervals configured, Hello 20 s, Dead 80 s, Retransmit 5 s
FSPF State is FULL
Neighbor Domain Id is 0x0c(12), Neighbor Interface index is 0x0f100000
Statistics counters :
  Number of packets received : LSU 8 LSA 8 Hello 118 Error packets 0
  Number of packets transmitted : LSU 8 LSA 8 Hello 119 Retransmitted LSU 0
  Number of times inactivity timer expired for the interface = 0
```

Default Settings

[Table 24-4](#) lists the default settings for FSPF features.

Table 24-4 **Default FSPF Settings**

| Parameters | Default |
|-----------------|--------------------------------------|
| FSPF | Enabled on all E ports and TE ports. |
| SPF computation | Dynamic. |
| SPF hold time | 0. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 24-4 **Default FSPF Settings (continued)**

| Parameters | Default |
|--|---|
| Backbone region | 0. |
| Acknowledgment interval (RxmtInterval) | 5 seconds. |
| Refresh time (LSRefreshTime) | 30 minutes. |
| Maximum age (MaxAge) | 60 minutes. |
| Hello interval | 20 seconds. |
| Dead interval | 80 seconds. |
| Distribution tree information | Derived from the principal switch (root node). |
| Routing table | FSPF stores up to 16 equal cost paths to a given destination. |
| Load balancing | Based on destination ID and source ID on different, equal cost paths. |
| In-order delivery | Disabled. |
| Drop latency | Disabled. |
| Static route cost | If the cost (metric) of the route is not specified, the default is 10. |
| Remote destination switch | If the remote destination switch is not specified, the default is direct. |
| Multicast routing | Uses the principal switch to compute the multicast tree. |



Configuring Intelligent Storage Services

Intelligent Storage Services are features supported on the Storage Services Module (SSM). You can use Intelligent Storage Services on Cisco MDS 9000 Family switches running the Cisco MDS SAN-OS Release 2.0(2b) or later software. Intelligent Storage Services supported in Cisco MDS SAN-OS Release 2.0(2b) include the following:

- Fibre Channel write acceleration
- SCSI flow statistics

Intelligent Storage Services supported in Cisco MDS SAN-OS Release 2.1(1a) include the following:

- SANTap
- Network-Accelerated Serverless Backup (NASB)

This chapter includes the following sections:

- [About SCSI Flow Services, page 25-2](#)
- [Configuring SCSI Flow Services, page 25-3](#)
- [About Fibre Channel Write Acceleration, page 25-5](#)
- [Enabling Fibre Channel Write Acceleration, page 25-5](#)
- [About SCSI Flow Statistics, page 25-6](#)
- [Enabling SCSI Flow Statistics, page 25-7](#)
- [Displaying SCSI Flow Services Information, page 25-7](#)
- [About SANTap, page 25-10](#)
- [Enabling SANTap, page 25-14](#)
- [Displaying SANTap Information, page 25-15](#)
- [About NASB, page 25-17](#)
- [Enabling NASB, page 25-19](#)
- [Displaying NASB Information, page 25-20](#)
- [Default Settings, page 25-21](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

About SCSI Flow Services

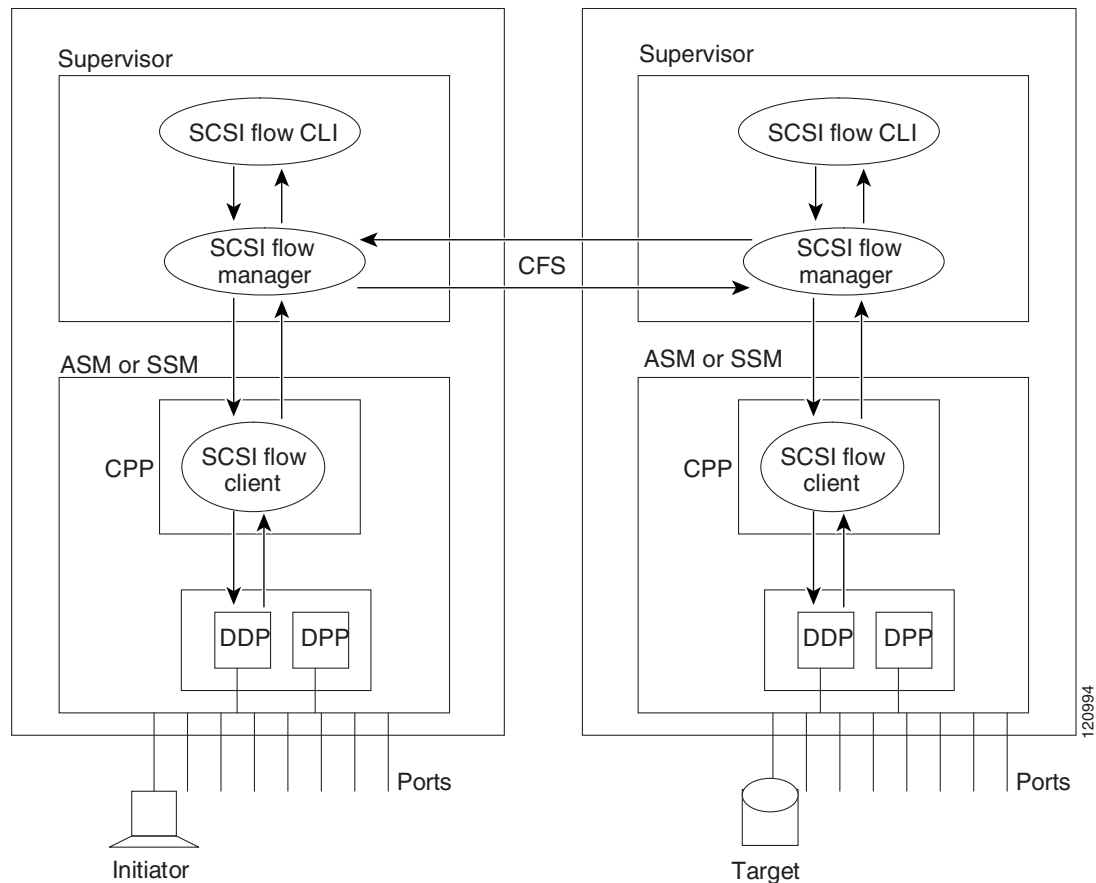
A SCSI initiator/target combination is a SCSI flow. SCSI Flow Services provide enhanced features for SCSI flows such as write acceleration and flow monitoring for statistics gathering on an SSM.

Functionally, the SCSI Flow Services functional architecture consists of the following components:

- SCSI flow manager (SFM) on the supervisor
- SCSI flow configuration CLI on the supervisor
- SCSI flow configuration client on the Control Path Processor (CPP) of an SSM
- SCSI flow feature set support on the Data Path Processor (DPP) of an SSM

Figure 25-1 shows an example of the SCSI Flow Services functional architecture.

Figure 25-1 SCSI Flow Services Functional Architecture



Note

For statistics monitoring, the target device is not required to be connected to an SSM.

Send documentation comments to mdsfeedback-doc@cisco.com.

SCSI Flow Manager

The SCSI flow manager (SFM) resides on a supervisor module and handles the configuration of SCSI flows, validating them and relaying configuration information to the appropriate SSM. It also handles any dynamic changes to the status of the SCSI flow due to external events. The SFM registers events resulting from operations, such as port up or down, VSAN suspension, and zoning that affects the SCSI flow status, and updates the flow status and configuration accordingly.

The SFM on the initiator communicates to its peer on the target side using Cisco Fabric Services (CFS). Peer communication allows the initiator SFM to validate target parameters and program information on the target side.

SCSI Flow Configuration Client

A SCSI flow configuration client (SFCC) resides on the CPP of the SSM. It receives flow configuration requests from the SFM, programs the DPP corresponding to the initiator and target port interfaces, and responds to the SFM with the status of the configuration request.

SCSI Flow Data Path Support

The DPP on the SSM examines all the messages between the initiator and target and provides SCSI flow features such as Fibre Channel write acceleration and statistics monitoring.

Configuring SCSI Flow Services

A SCSI flow specification consists of the following attributes:

- SCSI flow identifier
- VSAN identifier
- SCSI initiator port WWN
- SCSI target port WWN
- Flow feature set consisting of Fibre Channel write acceleration and statistics monitoring.

The SCSI flow specification is a distributed configuration because the SCSI initiator and the target might be physically connected to SSMs on two different switches located across the fabric. The configuration does not require information to identify either the switch name or the SSM slot location for either the initiator or the target. The manual SCSI flow configuration is performed only at the initiator side. This simplifies the configuration process. The initiator switch sends the configuration to the SFM on the target switch using CFS. No SCSI flow configuration is necessary on the target switch.

Send documentation comments to mdsfeedback-doc@cisco.com.

Enabling SCSI Flow Services

To enable SCSI Flow Services, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ssm enable feature scsi-flow module 2 | Enables SCSI Flow Services on the SSM in slot 2. |
| | switch(config)# no ssm enable feature scsi-flow module 2 | Disables SCSI Flow Services on the SSM in slot 2. The default is disabled. |
| | switch(config)# no ssm enable feature scsi-flow force module 2 | Forces the switch to disable of SCSI Flow Services on the SSM in slot 2. The default is disabled. |
| Step 3 | switch(config)# ssm enable feature scsi-flow interface fc 2/5 - 8 | Enables SCSI Flow Services on the interface 5 through 8 on the SSM in slot 2. Note Interfaces must be specified in multiples of four beginning at ports 1, 5, 9, 13, 17, 21, 25, and 29. |
| | switch(config)# no ssm enable feature scsi-flow interface fc 2/5 - 8 | Disables SCSI Flow Services on the interface 5 through 8 on the SSM in slot 2. The default is disabled. |
| | switch(config)# no ssm enable feature scsi-flow force interface fc 2/5 - 8 | Forces the switch to disable SCSI Flow Services on the interface 5 through 8 on the SSM in slot 2. |

Enabling SCSI Flow Configuration Distribution

To enable SCSI flow configuration distribution using CFS, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# scsi-flow distribute | Enables SCSI flow configuration distribution through CFS. The default is enabled. |
| | switch(config)# no scsi-flow distribute | Disables CFS distribution for SCSI flow configuration. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Configuring SCSI Flow Identifiers

A SCSI flow identifier is unique on a switch and is chosen by the user, like VSAN identifiers. To configure a SCSI flow identifier, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# scsi-flow flow-id 3 initiator-vsan 2 initiator-pwwn 21:00:00:e0:8b:07:5f:aa target-vsan 4 target-pwwn 2a:20:00:05:30:00:77:e0 | Configures SCSI flow identifier 3 using the pWWNs of the initiator and the target. The flow identifier range is 1 to 65535. |
| | switch(config)# no scsi-flow flow-id 3 initiator-vsan 2 | Removes a SCSI flow identifier 3. |

About Fibre Channel Write Acceleration

Fibre Channel write acceleration minimizes application latency or reduces transactions per second over long distances. For synchronous data replication, Fibre Channel write acceleration increases the distance of replication or reduces effective latency to improve performance. To take advantage of this feature, both the initiator and target devices must be directly attached to an SSM.

The Fibre Channel write acceleration feature also allows the configuration of the buffer count. You can change the number of 2 KB buffers reserved on the target side DPP for a SCSI flow.

You can estimate the number of buffers to configure using the following formula:

(Number of concurrent SCSI writes * size of SCSI writes in bytes) / FCP data frame size in bytes

For example, for HDS TrueCopy between HDS 9970s, which use 1KB FCP data frames, and you perform an initial sync for a 16-LUN TrueCopy group with 15 tracks, or 768 KB per LUN, the approximate number of write buffers required would be 16*(768*1024)/1024 or 12248 buffers.



Note

The Fibre Channel write acceleration feature requires the Enterprise Package license installed on both the initiator and target switches.

Enabling Fibre Channel Write Acceleration

To enable Fibre Channel write acceleration, and optionally modify the number of write acceleration buffers, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# scsi-flow flow-id 3 write-acceleration | Enables Fibre Channel write acceleration for SCSI flow identifier 3. |
| | switch(config)# no scsi-flow flow-id 3 write-acceleration | Disables SCSI flow write acceleration for SCSI flow identifier 3. The default is disabled. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|--|--|
| Step 3 | <code>switch(config)# scsi-flow flow-id 3 write-acceleration buffer 2048</code> | Enables Fibre Channel write acceleration for SCSI flow identifier 3 and sets the number of buffers to 2048. The range is 1 to 40000. |
| | <code>switch(config)# no scsi-flow flow-id 3 write-acceleration buffer 1024</code> | Reverts to the default number of write acceleration buffers. The default is 1024. |

About SCSI Flow Statistics

The statistics that can be collected for SCSI flows include the following:

- SCSI reads
 - Number of I/O s
 - Number of I/O blocks
 - Maximum I/O blocks
 - Minimum I/O response time
 - Maximum I/O response time
- SCSI writes
 - Number of I/Os
 - Number of I/O blocks
 - Maximum I/O blocks
 - Minimum I/O response time
 - Maximum I/O response time
- Other SCSI commands (not read or write)
 - Test unit ready
 - Report LUN
 - Inquiry
 - Read capacity
 - Mode sense
 - Request sense
- Errors
 - Number of timeouts
 - Number of I/O failures
 - Number of various SCSI status events
 - Number of various SCSI sense key errors or events

To take advantage of this feature, only the initiator must be directly attached to an SSM.



Note

The SCSI flow statistics feature requires the Enterprise Package license installed only on the initiator switches.

Send documentation comments to mdsfeedback-doc@cisco.com.

Enabling SCSI Flow Statistics

To enable SCSI flow statistics monitoring, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# scsi-flow flow-id 3 statistics | Enables statistics monitoring on SCSI flow identifier 3. |
| | switch(config)# no scsi-flow flow-id 3 statistics | Disables statistics monitoring on SCSI flow identifier 3. The default is disabled. |

Clearing SCSI Flow Statistics

Use the **clear device-name statistics flow-id** command to clear SCSI flow statistics (for debugging purposes):

```
switch# clear scsi-flow statistics flow-id 3
```

Displaying SCSI Flow Services Information

Use the **show scsi-flow** command to display information about SCSI Flow Services (see [Example 25-1](#) to [Example 25-5](#)).

Example 25-1 Displays Applications Provisioned on an SSM

```
switch# show ssm provisioning
Module   Ports      Application      Provisioning Status
-----
4        1-32      scsi-flow        success
```

Example 25-2 Displays SCSI Flow Services Configuration for All SCSI Flow Identifiers

```
switch# show scsi-flow
Flow Id: 3
Initiator VSAN: 101
Initiator WWN: 21:00:00:e0:8b:05:76:28
Target VSAN: 102
Target WWN: 21:00:00:20:37:38:7f:7d
Target LUN: ALL LUNs
Flow Verification Status:
-----
Initiator Verification Status: success
Target Verification Status: success
Initiator Linecard Status: success
Target Linecard Status: success
Feature Status:
-----
Write-Acceleration enabled
Write-Acceleration Buffers: 1024
Configuration Status: success
Statistics enabled
Configuration Status: success
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
Flow Id: 4
  Initiator VSAN: 101
  Initiator WWN: 21:00:00:e0:8b:05:76:28
  Target VSAN: 102
  Target WWN: 21:00:00:20:37:38:a7:89
  Target LUN: ALL LUNs
  Flow Verification Status:
  -----
    Initiator Verification Status:  success
    Target Verification Status:      success
    Initiator Linecard Status:      success
    Target Linecard Status:         success
  Feature Status:
  -----
    Write-Acceleration enabled
    Write-Acceleration Buffers: 1024
    Configuration Status:  success
```

Example 25-3 Displays SCSI Flow Services Configuration for a Specific SCSI Flow Identifier

```
switch# show scsi-flow flow-id 3
Flow Id: 3
  Initiator VSAN: 101
  Initiator WWN: 21:00:00:e0:8b:05:76:28
  Target VSAN: 102
  Target WWN: 21:00:00:20:37:38:7f:7d
  Target LUN: ALL LUNs
  Flow Verification Status:
  -----
    Initiator Verification Status:  success
    Target Verification Status:      success
    Initiator Linecard Status:      success
    Target Linecard Status:         success
  Feature Status:
  -----
    Write-Acceleration enabled
    Write-Acceleration Buffers: 1024
    Configuration Status:  success
    Statistics enabled
    Configuration Status:  success
```

Example 25-4 Displays SCSI Flow Services Statistics for All SCSI Flow Identifiers

```
switch# show scsi-flow statistics

Stats for flow-id 4 LUN=0x0000
-----
Read Stats
  I/O Total count=2
  I/O Timeout count=0
  I/O Total block count=4
  I/O Max block count=2
  I/O Min response time=5247 usec
  I/O Max response time=10160 usec
  I/O Active Count=0
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
Write Stats
I/O Total count=199935
I/O Timeout count=0
I/O Total block count=12795840
I/O Max block count=64
I/O Min response time=492 usec
I/O Max response time=10056529 usec
I/O Active Count=16

Non Read-Write Stats
Test Unit Ready=4
Report LUN=38
Inquiry=50
Read Capacity=3
Mode Sense=0
Request Sense=0

Total Stats
Rx Frame Count=3792063
Rx Frame Byte Count=6549984752
Tx Frame Count=3792063
Tx Frame Byte Count=6549984752

Error Stats
SCSI Status Busy=0
SCSI Status Reservation Conflict=0
SCSI Status Task Set Full=0
SCSI Status ACA Active=0
Sense Key Not Ready=0
Sense Key Medium Error=0
Sense Key Hardware Error=0
Sense Key Illegal Request=0
Sense Key Unit Attention=28
Sense Key Data Protect=0
Sense Key Blank Check=0
Sense Key Copy Aborted=0
Sense Key Aborted Command=0
Sense Key Volume Overflow=0
Sense Key Miscompare=0
```

Example 25-5 Displays SCSI Flow Services Statistics for a Specific SCSI Flow Identifier

```
switch# show scsi-flow statistics flow-id 4
```

```
Stats for flow-id 4 LUN=0x0000
-----
Read Stats
I/O Total count=2
I/O Timeout count=0
I/O Total block count=4
I/O Max block count=2
I/O Min response time=5247 usec
I/O Max response time=10160 usec
I/O Active Count=0

Write Stats
I/O Total count=199935
I/O Timeout count=0
I/O Total block count=12795840
I/O Max block count=64
I/O Min response time=492 usec
I/O Max response time=10056529 usec
I/O Active Count=16
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```

Non Read-Write Stats
Test Unit Ready=4
Report LUN=38
Inquiry=50
Read Capacity=3
Mode Sense=0
Request Sense=0

Total Stats
Rx Frame Count=3792063
Rx Frame Byte Count=6549984752
Tx Frame Count=3792063
Tx Frame Byte Count=6549984752

Error Stats
SCSI Status Busy=0
SCSI Status Reservation Conflict=0
SCSI Status Task Set Full=0
SCSI Status ACA Active=0
Sense Key Not Ready=0
Sense Key Medium Error=0
Sense Key Hardware Error=0
Sense Key Illegal Request=0
Sense Key Unit Attention=28
Sense Key Data Protect=0
Sense Key Blank Check=0
Sense Key Copy Aborted=0
Sense Key Aborted Command=0
Sense Key Volume Overflow=0
Sense Key Miscompare=0


```

About SANTap

The SANTap feature allows third party data storage applications, such as long distance replication and continuous backup, to be integrated into the SAN. The protocol-based interface that is offered by SANTap allows easy and rapid integration of the data storage service application because it delivers a loose coupling between the application and an SSM, thereby reducing the effort needed to integrate applications with the core services being offered by the SSM. See [Figure 25-2](#).

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 25-2 *Integrating Third-Party Storage Applications in a SAN*




Send documentation comments to mdsfeedback-doc@cisco.com.

SANTap operates in three modes:

- Transparent mode

Transparent mode eliminates the need for any reconfiguration of either the host or target when introducing SANTap based applications. This mode of operation requires that either the host initiator or target be directly connected to an SSM. See [Figure 25-3](#).

Figure 25-3 *SANTap Transparent Mode Example*




Send documentation comments to mdsfeedback-doc@cisco.com.

- Proxy mode-1

Proxy mode-1 assigns Cisco-specific WWNs to the virtual initiators (VIs) and digital virtual targets (DVTs). The benefit of this mode is that it eliminates the requirement of transparent mode that a host initiator or a target be connected directly to an SSM. In proxy mode-1, the SSM can be anywhere in the SAN. However, this mode requires reconfiguration of legacy applications. See [Figure 25-4](#).

Figure 25-4 SANTap Proxy Mode-1 Example




130077

Send documentation comments to mdsfeedback-doc@cisco.com.

- Proxy mode-2

Proxy mode-2 includes the benefits of transparent mode and proxy mode-1 but does not have the limitations of those modes. However, it does require that the administrator partition the SAN using VSANs. The host initiator and the DVT are in one VSAN while the VI and the target are in another VSAN. See [Figure 25-5](#).

Figure 25-5 SANTap Proxy Mode-2 Example



Enabling SANTap

SANTap can be enabled on an entire SSM or it can be enabled on a group of four ports on an SSM. To enable the SANTap feature, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ssm enable feature santap module 4 | Enables the SANTap application on the entire SSM. |
| | switch(config)# no ssm enable feature santap module 4 | Disables the SANTap application on the entire SSM in slot 4. |
| | switch(config)# no ssm enable feature santap force module 4 | Forces the switch to disable the SANTap application on the entire SSM in slot 4. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|--|--|
| Step 3 | <code>switch(config)# ssm enable feature santap interface fc 4/1 - 4</code> | Enables the SANTap application on ports 1 through 4 on the SSM. Note Interfaces must be specified in multiples of four beginning at ports 1, 5, 9, 13, 17, 21, 25, and 29. |
| | <code>switch(config)# no ssm enable feature santap interface fc 4/1 - 4</code> | Disables the SANTap application on ports 1 through 4 on the SSM in slot 4. |
| | <code>switch(config)# no ssm enable feature santap force interface fc 4/1 - 4</code> | Forces the switch to disable the SANTap application on ports 1 through 4 on the SSM in slot 4. |
| Step 4 | <code>switch(config)# santap module 4 appl-vsan 10</code> | Enables SANTap on the SSM in slot 4 and on VSAN 10. |
| | <code>switch(config)# no santap module 4 appl-vsan 10</code> | Disables SANTap. |

Displaying SANTap Information

Use the `show santap module` command to display information about SANTap (see [Example 25-6](#) to [Example 25-13](#)).

Example 25-6 Displays SANTap Control Virtual Terminal Information

```
switch# show santap module 2 cvt

CVT Information :
  cvt pwwn      = 25:3c:00:05:30:00:22:25
  cvt nwwn      = 25:3d:00:05:30:00:22:25
  cvt id        = 1
  cvt xmap_id   = 2
  cvt vsan      = 10
```

Example 25-7 Displays SANTap Data Virtual Terminal Information

```
switch# show santap module 2 dvt

DVT Information :
  dvt pwwn      = 22:00:00:20:37:88:20:ef
  dvt nwwn      = 20:00:00:20:37:88:20:ef
  dvt id        = 3
  dvt mode      = 3
  dvt vsan      = 3
  dvt fp_port   = 0
  dvt if_index  = 0x1080000
  dvt name      = MYDVT
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 25-8 Displays SANTap Data Virtual Terminal LUN Information

```
switch# show santap module 2 dvtlun

DVT LUN Information :
  dvt pwwn      = 22:00:00:20:37:88:20:ef
  dvt lun       = 0x0
  xmap id       = 8
  dvt id        = 3
  dvt mode      = 0
  dvt vsan      = 3
  tgt pwwn      = 22:00:00:20:37:88:20:ef
  tgt lun       = 0x0
  tgt vsan      = 1
```

Example 25-9 Displays SANTap Session Information

```
switch# show santap module 2 session

Session Information :
  session id    = 1
  host pwwn     = 21:00:00:e0:8b:07:61:aa
  dvt pwwn      = 22:00:00:20:37:88:20:ef
  dvt lun       = 0x0
  tgt pwwn      = 00:00:00:00:00:00:00:00
  tgt lun       = 0x0
  adt pwwn      = 77:77:77:77:77:77:77:77
  adt lun       = 0x0
  num ranges    = 0
  dvt id        = 0
  vdisk id      = 0
  session state = 0
  mrl requested = 1
  pwl requested = 1
  iol requested = 0
```

Example 25-10 Displays SANTap Appliance Virtual Terminal Information

```
switch# show santap module 2 avt

AVT Information :
  avt pwwn      = 2a:4b:00:05:30:00:22:25
  avt nwwn      = 2a:60:00:05:30:00:22:25
  avt id        = 12
  avt vsan      = 4
  avt if_index  = 0x1080000
  hi pwwn      = 21:00:00:e0:8b:07:61:aa
  tgt pwwn      = 22:00:00:20:37:88:20:ef
  tgt vsan      = 1
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 25-11 Displays SANTap Appliance Virtual Terminal LUN Information

```
switch# show santap module 2 avtlun

AVT LUN Information :
  avt pwwn      = 2a:4b:00:05:30:00:22:25
  avt lun       = 0x0
  xmap id       = 16
  avt id        = 12
  tgt lun       = 0x0
```

Example 25-12 Displays SANTap Remote Virtual Terminal Information

```
switch# show santap mod 2 rvt

RVT Information :
  rvt pwwn      = 2a:61:00:05:30:00:22:25
  rvt nwwn      = 2a:62:00:05:30:00:22:25
  rvt id        = 17
  rvt vsan      = 4
  rvt if_index  = 0x1080000
```

Example 25-13 Displays SANTap Remote Virtual Terminal LUN Information

```
switch# show santap mod 2 rvtlun

RVT LUN Information :
  rvt pwwn      = 2a:61:00:05:30:00:22:25
  rvt lun       = 0x0
  xmap id       = 22
  rvt id        = 17
  app pwwn      = 22:00:00:20:37:39:b1:00
  app lun       = 0x0
  app vsan      = 1
```

About NASB

As of Cisco MDS SAN-OS Release 2.1(1a), the SSMs support Network-Accelerated Serverless Backup (NASB).

Data movement in the fabric uses considerable processor cycles, which can cause client applications to slow down noticeably. Offloading data movement operations to a media server allows the client applications to run normally even during a backup operation. Media servers can further offload the data movement operation to NASB devices, which allows the media server to focus on the coordination functions needed to complete the backup.

Most backups performed today are server-free. In server-free backups, the application server is not involved in moving the data. The data can be moved by either a media server or an NASB device.

When the media server is the data mover, it moves the data between the disks and the tapes. The backup application runs on both the client device and the media server. However, the backup application in the client device performs minimal tasks for the backup operation.

Send documentation comments to mdsfeedback-doc@cisco.com.


The media server performs the following backup operations:

- Manages disks as well as one or more tape backup devices.
- Contacts the client devices to retrieve the list of logical blocks that need to be backed up.
- Performs data movement from disk to tape media based on the logical block list provided by the client device.

The backup application in the client device maps the data to be backed up and creates the logical block list associated with the data. The movement of data from the physical disks to the backup device (tape) is not performed by the client device. This reduces substantial load on the client device.

An example configuration is shown in [Figure 25-6](#). The media server moves the data directly between the storage disks and the tape devices during backups.


Figure 25-6 Example Configuration with Media Server as Data Mover



When the NASB is the data mover, it moves the data between the disks and the tapes. The NASB device is a SCSI target device capable of handling SCSI Extended Copy (XCOPY) commands as well as a SCSI initiator device capable of issuing READ/WRITE commands to disks and other backup media, such as tapes. See [Figure 25-7](#).

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 25-7 Example Configuration with NASB Device as Data Mover



The task of managing and preparing the source and destination targets is performed by the media server. For example, if the destination is a tape library, the media server issues commands to load and unload the correct tape and position of the tape write head at the correct offset within the tape.

Enabling NASB

Network-accelerated Serverless Backup (NASB) can be enabled on an entire SSM or it can be enabled on one or more groups of four ports on an SSM. To enable the NASB feature, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ssm enable feature nasb module 4 | Enables the NASB application on the entire SSM in slot 4. |
| | switch(config)# no ssm enable feature nasb module 4 | Disables the NASB application on the entire SSM in slot 4. |
| | switch(config)# no ssm enable feature nasb force module 4 | Forces the switch to disable the NASB application on the entire SSM in slot 4. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|--|--|
| Step 3 | <code>switch(config)# ssm enable feature nasb interface fc 4/1 - 4</code> | Enables the NASB application on ports 1 through 4 on the SSM in slot 4. Note Interfaces must be specified in multiples of four beginning at ports 1, 5, 9, 13, 17, 21, 25, and 29. |
| | <code>switch(config)# no ssm enable feature nasb interface fc 4/1 - 4</code> | Disables the NASB application on ports 1 through 4 on the SSM in slot 4. |
| | <code>switch(config)# no ssm enable feature nasb force interface fc 4/1 - 4</code> | Forces the switch to disable the NASB application on ports 1 through 4 on the SSM in slot 4. |
| Step 4 | <code>switch(config)# nasb module 4 vsan 10</code> | Enables NASB on the SSM in slot 4 and on VSAN 10 for a single target LUN. |
| | <code>switch(config)# nasb module 4 vsan 10 multiple</code> | Enables NASB on the SSM in slot 4 and on VSAN 10 for multiple target LUNs. |
| | <code>switch(config)# no nasb module 4 vsan 10</code> | Disables NASB. |

Displaying NASB Information

Use the **show nasb** command to display information about NASB (see [Example 25-14](#) to [Example 25-17](#)).

Example 25-14 Displays NASB Information

```
switch# show nasb
NASB:module 3 vsan 1:DPP-1, VT-nWWN=22f90005300036a2, pWWN=22fa0005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-2, VT-nWWN=22fb0005300036a2, pWWN=22fc0005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-3, VT-nWWN=22fd0005300036a2, pWWN=22fe0005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-4, VT-nWWN=22ff0005300036a2, pWWN=26000005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-5, VT-nWWN=26010005300036a2, pWWN=26020005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-6, VT-nWWN=26030005300036a2, pWWN=26040005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-7, VT-nWWN=26050005300036a2, pWWN=26060005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-8, VT-nWWN=26070005300036a2, pWWN=26080005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-1, VT-nWWN=26090005300036a2, pWWN=260a0005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-2, VT-nWWN=260b0005300036a2, pWWN=260c0005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-3, VT-nWWN=260d0005300036a2, pWWN=260e0005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-4, VT-nWWN=260f0005300036a2, pWWN=26100005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-5, VT-nWWN=26110005300036a2, pWWN=26120005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-6, VT-nWWN=26130005300036a2, pWWN=26140005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-7, VT-nWWN=26150005300036a2, pWWN=26160005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-8, VT-nWWN=26170005300036a2, pWWN=26180005300036a2 (provisioned)
```

Example 25-15 Displays NASB Information for a Specific Module

```
switch# show nasb module 3
NASB:module 3 vsan 1:DPP-1, VT-nWWN=22f90005300036a2, pWWN=22fa0005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-2, VT-nWWN=22fb0005300036a2, pWWN=22fc0005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-3, VT-nWWN=22fd0005300036a2, pWWN=22fe0005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-4, VT-nWWN=22ff0005300036a2, pWWN=26000005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-5, VT-nWWN=26010005300036a2, pWWN=26020005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-6, VT-nWWN=26030005300036a2, pWWN=26040005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-7, VT-nWWN=26050005300036a2, pWWN=26060005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-8, VT-nWWN=26070005300036a2, pWWN=26080005300036a2 (provisioned)
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
NASB:module 3 vsan 2:DPP-1, VT-nWWN=26090005300036a2, pWWN=260a0005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-2, VT-nWWN=260b0005300036a2, pWWN=260c0005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-3, VT-nWWN=260d0005300036a2, pWWN=260e0005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-4, VT-nWWN=260f0005300036a2, pWWN=26100005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-5, VT-nWWN=26110005300036a2, pWWN=26120005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-6, VT-nWWN=26130005300036a2, pWWN=26140005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-7, VT-nWWN=26150005300036a2, pWWN=26160005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-8, VT-nWWN=26170005300036a2, pWWN=26180005300036a2 (provisioned)
```

Example 25-16 Displays NASB Information for a Specific Module for a VSAN

```
switch# show nasb module 3 vsan 2
NASB:module 3 vsan 2:DPP-1, VT-nWWN=26090005300036a2, pWWN=260a0005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-2, VT-nWWN=260b0005300036a2, pWWN=260c0005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-3, VT-nWWN=260d0005300036a2, pWWN=260e0005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-4, VT-nWWN=260f0005300036a2, pWWN=26100005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-5, VT-nWWN=26110005300036a2, pWWN=26120005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-6, VT-nWWN=26130005300036a2, pWWN=26140005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-7, VT-nWWN=26150005300036a2, pWWN=26160005300036a2 (provisioned)
NASB:module 3 vsan 2:DPP-8, VT-nWWN=26170005300036a2, pWWN=26180005300036a2 (provisioned)
```

Example 25-17 Displays NASB Information for a Specific VSAN

```
switch# show nasb vsan 1
NASB:module 3 vsan 1:DPP-1, VT-nWWN=22f90005300036a2, pWWN=22fa0005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-2, VT-nWWN=22fb0005300036a2, pWWN=22fc0005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-3, VT-nWWN=22fd0005300036a2, pWWN=22fe0005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-4, VT-nWWN=22ff0005300036a2, pWWN=26000005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-5, VT-nWWN=26010005300036a2, pWWN=26020005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-6, VT-nWWN=26030005300036a2, pWWN=26040005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-7, VT-nWWN=26050005300036a2, pWWN=26060005300036a2 (provisioned)
NASB:module 3 vsan 1:DPP-8, VT-nWWN=26070005300036a2, pWWN=26080005300036a2 (provisioned)
```

Default Settings

Table 25-1 lists the default settings for Intelligent Storage Services parameters.

Table 25-1 Default Intelligent Storage Services Parameters

| Parameters | Default |
|--|----------|
| SCSI Flow Services | Disabled |
| SCSI Flow Services distribution | Enabled |
| Fibre Channel write acceleration | Disabled |
| Fibre Channel write acceleration buffers | 1024 |
| SCSI Flow Services statistics | Disabled |
| SANTap feature | Disabled |
| NASB feature | Disabled |

Send documentation comments to mdsfeedback-doc@cisco.com.



Configuring IP Services

Cisco MDS 9000 Family switches can route IP traffic between Ethernet and Fibre Channel interfaces. The IP static routing feature is used to route traffic between VSANs. To do so, each VSAN must be in a different IP subnetwork. Each Cisco MDS 9000 Family switch provides the following services for network management systems (NMS):

- IP forwarding on the out-of-band Ethernet interface (mgmt0) on the front panel of the supervisor modules.
- IP forwarding or in-band Fibre Channel interface using the IP over Fibre Channel (IPFC) function—IPFC specifies how IP frames can be transported over Fibre Channel using encapsulation techniques. IP frames are encapsulated into Fibre Channel frames so NMS information can cross the Fibre Channel network without using an overlay Ethernet network.
- IP routing (default routing and static routing)—If your configuration does not need an external router, you can configure a default route using static routing.

Switches are compliant with RFC 2338 standards for Virtual Router Redundancy Protocol (VRRP) features. VRRP is a restartable application that provides a redundant, alternate path to the gateway switch.

This chapter includes the following sections:

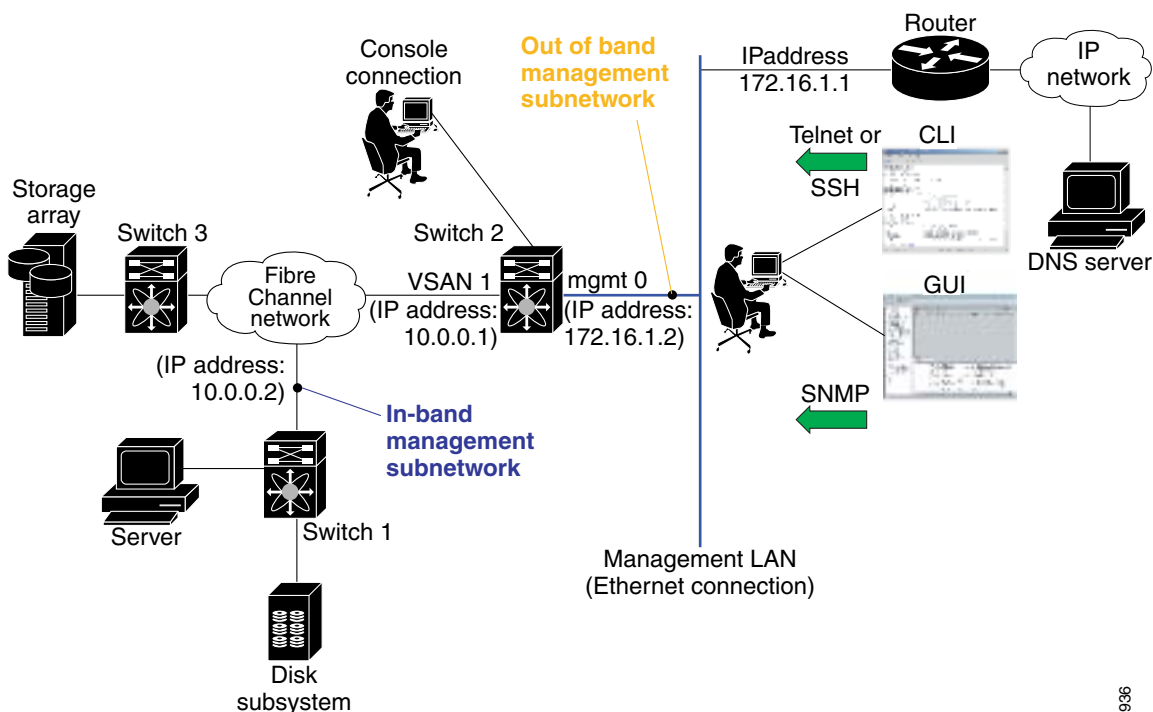
- [Traffic Management Services, page 26-2](#)
- [Management Interface Configuration, page 26-2](#)
- [Default Gateway Configuration, page 26-3](#)
- [Default Network Configuration, page 26-3](#)
- [IP Access Control Lists, page 26-5](#)
- [IPFC Configuration, page 26-12](#)
- [Configuring IP Static Routes, page 26-13](#)
- [Displaying IP Interface Information, page 26-14](#)
- [Overlay VSAN Configuration, page 26-15](#)
- [Multiple VSAN Configuration, page 26-17](#)
- [The Virtual Router Redundancy Protocol, page 26-19](#)
- [DNS Server Configuration, page 26-25](#)
- [Default Settings, page 26-26](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

Traffic Management Services

In-band options are compliant with and use the RFC 2625 standards. An NMS host running IP protocol over a FC interface can access the switch using the IPFC functionality. If the NMS does not have a Fibre Channel HBA, in-band management can still be performed using one of the switches as an access point to the fabric (see Figure 26-1).

Figure 26-1 Management Access to Switches



79936

Management Interface Configuration

On director class switches, a single IP address is used to manage the switch. The active supervisor module's management (mgmt0) interface uses this IP address. The mgmt0 interface on the standby supervisor module remains in an inactive state and cannot be accessed until a switchover happens. After a switchover, the mgmt0 interface on the standby supervisor module becomes active and assumes the same IP address as the previously-active supervisor module.

The management interface on the switch allows multiple simultaneous Telnet or SNMP sessions. You can remotely configure the switch through the management interface, but first you must configure some IP parameters (IP address, subnet mask) so that the switch is reachable. You can manually configure the management interface from the CLI.



Note

Before you begin to configure the management interface manually, obtain the switch's IP address and IP subnet mask. Also make sure the console cable is connected to the console port.

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure the mgmt0 Ethernet interface, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# interface mgmt0 switch(config-if)# | Enters the interface configuration mode on the management Ethernet interface (mgmt0). |
| Step 3 | switch(config-if)# ip address 1.1.1.1 255.255.255.0 | Enters the IP address (1.1.1.1) and IP subnet mask (255.255.255.0) for the management interface. |
| Step 4 | switch(config-if)# no shutdown | Enables the interface. |

Default Gateway Configuration

The default gateway IP address should be configured along with the IP static routing commands (IP default network, destination prefix, and destination mask, and next hop address).



Tip

If you configure the static route IP forwarding and the default-network details, these IP addresses will be used regardless of the default-gateway being enabled or disabled. If these IP addresses are configured but not available, the switch will fall back to using the default gateway IP address, if you have configured it. Be sure to configure IP addresses for all entries in the switch.

See the [“Initial Setup Routine” section on page 4-2](#) for more information on configuring the IP addresses for all entries in the switch.

Use the **IP default-gateway** command to configure the IP address for a switch’s default gateway and the **show ip route** command to verify that the IP address for the default gateway is configured.

To configure default gateways, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# ip default-gateway 1.12.11.1 | Configures the IP address for the default gateway. |

Default Network Configuration

If you assign the IP default network address, the switch considers routes to that network as the last resort. If the IP default network address is not available, the switch uses the IP default gateway address. For every network configured with the IP default network address, the switch flags that route as a candidate default route, if the route is available.



Tip


If you configure the static route IP forwarding and the default network details, these IP addresses will be used regardless of the default gateway being enabled or disabled. If these IP address are configured and not available, the switch will fall back to using the default gateway IP address, if you have configured it. Be sure to configure IP addresses for all entries in the switch.

See the [“Initial Setup Routine” section on page 4-2](#) for more information on configuring the IP addresses for all entries in the switch.

Send documentation comments to mdsfeedback-doc@cisco.com.

When the Ethernet interface is configured, the switch should point to the gateway router for the IP network. The host accesses the gateway using a gateway switch. This gateway switch is configured as the default gateway. The other switches in the fabric that are connected to the same VSAN as the gateway switch can also be connected through the gateway switch. Every interface connected to this VSAN should be configured with the VSAN IP address of the gateway switch (see [Figure 26-2](#)).

Figure 26-2 Overlay VSAN Functionality



In [Figure 26-2](#), switch A has the IP address 1.1.2.11.1, switch B has the IP address 1.1.2.11.2, switch C has the IP address 1.1.2.11.3, and switch D has the IP address 1.1.2.11.4. Switch A is the gateway switch with the Ethernet connection. The NMS uses the IP address 1.1.1.10 to connect to the gateway switch. Frames forwarded to any switch in the overlaid VSAN 1 are routed through the gateway switch. Configuring the gateway switch's IP address, 1.1.2.11.1, in the other switches enable the gateway switch to forward the frame to the intended destination. Similarly, if a non-gateway switch in the VSAN forwards a frame to the Ethernet world, the frame is routed through the gateway switch.

When forwarding is disabled (default), IP frames are not sent from one interface to another. In these cases, the software performs local IP routing between two switches using the in-band option for Fibre Channel traffic and the mgmt0 option for Ethernet traffic.

When a VSAN is created, a VSAN interface is not created automatically. You need to specifically create the interface (see the [“Configuring VSAN Interfaces”](#) section on page 12-20).

Unlike the **ip default-gateway** command, use the **ip default-network** command when IP routing is enabled on the switch. Use the **show ip route** command to verify if the IP address for the default gateway is configured.

To configure default networks, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# ip default-network 190.10.1.0 | Configures the IP address for the default network (190.10.1.0). |
| | switch(config)# ip route 10.0.0.0 255.0.0.0 131.108.3.4 switch(config)# ip default-network 10.0.0.0 | Defines a static route to network 10.0.0.0 as the static default route. |

Send documentation comments to mdsfeedback-doc@cisco.com.

IP Access Control Lists

IP Access Control Lists (IP-ACLs) provide basic network security to all switches in the Cisco MDS 9000 Family. IP-ACLs restrict IP-related traffic based on the configured IP filters. A filter contains the rules to match an IP packet, and if the packet matches, the rule also stipulates if the packet should be permitted or denied.

Each switch in the Cisco MDS 9000 Family can have a maximum of 64 IP-ACLs, each IP-ACL can have a maximum of 256 filters.

IP-ACL Configuration Guidelines

Follow these guidelines when configuring IP-ACLs in any switch or director in the Cisco MDS 9000 Family:

- In Cisco SAN-OS Release 1.3 and earlier, you could only apply IP-ACLs to VSAN interfaces and the management interface. As of Cisco SAN-OS Release 2.0(1b), you can also apply IP-ACLs to Gigabit Ethernet interfaces (IPS modules) and Ethernet PortChannel interfaces.



Tip

If IP-ACLs are already configured in a Gigabit Ethernet interface, you cannot add this interface to a Ethernet PortChannel group. Refer to the [“Gigabit Ethernet IP-ACL Guidelines” section on page 28-10](#) for guidelines on configuring IP ACLs.



Caution

Do not apply IP-ACLs to only one member of a PortChannel group. Apply IP-ACLs to the entire channel group.

- Configure the order of conditions accurately. As the IP-ACL filters are sequentially applied to the IP flows, only the first match determines the action taken. Subsequent matches are not considered. Be sure to configure the most important condition first. If no conditions match, the software drops the packet.

Filter Contents

An IP filter contains rules for matching an IP packet based on the protocol, address, port, ICMP type, and type of service (TOS).

Protocol Information

The protocol information is required in each filter. It identifies the name or number of an IP protocol. You can specify the IP protocol in one of two ways:

- Specify an integer ranging from 0 to 255. This number represents the IP protocol.
- Specify the name of a protocol including, but not restricted to, Internet Protocol (IP, keyword **ip**), Transmission Control Protocol (TCP, keyword **tcp**), User Datagram Protocol (UDP, keyword **udp**), and Internet Control Message Protocol (ICMP, keyword **icmp**).



Note

When configuring IP-ACLs on Gigabit Ethernet interfaces, only use the TCP or ICMP options.

Send documentation comments to mdsfeedback-doc@cisco.com.

Address Information

The address information is required in each filter. It identifies the following details:

- Source: the address of the network or host from which the packet is being sent.
- Source-wildcard: the wildcard bits applied to the source.
- Destination: the number of the network or host to which the packet is being sent.
- Destination-wildcard: the wildcard bits applied to the destination.

Specify the source and source-wildcard or the destination and destination-wildcard in one of two ways:

- Using the 32-bit quantity in four-part, dotted decimal format (10.1.1.2/0.0.0.0 is the same as host 10.1.1.2).
 - Each wildcard bit set to zero indicates that the corresponding bit position in the packet's ip address must exactly match the bit value in the corresponding bit position in the source.
 - Each wildcard bit set to one indicates that both a zero bit and a one bit in the corresponding position of the packet's ip address will be considered a match to this access list entry. Place ones in the bit positions you want to ignore. For example, 0.0.255.255 to require an exact match of only the first 16 bits of the source. Wildcard bits set to one do not need to be contiguous in the source-wildcard. For example, a source-wildcard of 0.255.0.64 would be valid.
- Using the **any** option as an abbreviation for a source and source-wildcard or destination and destination-wildcard (0.0.0.0/255.255.255.255)

Port Information

The port information is optional. To compare the source and destination ports, use the **eq** (equal) option, the **gt** (greater than) option, the **lt** (less than) option, or the **range** (range of ports) option. You can specify the port information in one of two ways:

- Specify the number of the port. Port numbers range from 0 to 65535. [Table 26-1](#) displays the port numbers recognized by the Cisco SAN-OS software for associated TCP and UDP ports.
- Specify the name of a TCP or UDP port as follows:
 - TCP port names can only be used when filtering TCP.
 - UDP port names can only be used when filtering UDP.

Table 26-1 TCP and UDP Port Numbers

| Protocol | Port | Number |
|----------|-----------------------|--------------|
| UDP | dns | 53 |
| | tftp | 69 |
| | ntp | 123 |
| | radius accounting | 1646 or 1813 |
| | radius authentication | 1645 or 1812 |
| | snmp | 161 |
| | snmp-trap | 162 |
| | syslog | 514 |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 26-1 TCP and UDP Port Numbers (continued)

| Protocol | Port | Number |
|---|------------|--------|
| TCP | ftp | 20 |
| Note If the TCP connection is already established, use the established option to find matches. A match occurs if the TCP datagram has the ACK, FIN, PSH, RST, or URG control bit set. | ftp-data | 21 |
| | ssh | 22 |
| | telnet | 23 |
| | smtp | 25 |
| | tasacs-ds | 65 |
| | www | 80 |
| | sftp | 115 |
| | http | 143 |
| | wbem-http | 5988 |
| | wbem-https | 5989 |

ICMP Information

IP packets can be filtered based on the following optional ICMP conditions:

- The icmp-type: ICMP message type. The type is a number from 0 to 255.
- The icmp-code: ICMP message code. The code is a number from 0 to 255.

Table 26-2 displays the value for each ICMP type.

Table 26-2 ICMP Type Value

| ICMP Type ¹ | Code |
|-------------------------|------|
| echo | 8 |
| echo-reply | 0 |
| destination unreachable | 3 |
| traceroute | 30 |
| time exceeded | 11 |

1. ICMP redirect packets are always rejected.

TOS Information

IP packets can be filtered based on the following optional TOS conditions:

- The TOS level, as specified by a number from 0 to 15
- The TOS name: max-reliability, max-throughput, min-delay, min-monetary-cost, and normal

Send documentation comments to mdsfeedback-doc@cisco.com.

IP-ACL -Creation

Traffic coming into the switch is compared to IP-ACL filters based on the order that the filters occur in the switch. New filters are added to the end of the IP-ACL. The switch keeps looking until it has a match. If no matches are found when the switch reaches the end of the filter, the traffic is denied. For this reason, you should have the frequently hit filters at the top of the filter. There is an *implied deny* for traffic that is not permitted. A single-entry IP-ACL with only one **deny** entry has the effect of denying all traffic.

To configure an IP-ACL, you must complete the following tasks:

1. Create an IP-ACL by specifying a filter name and one or more access condition(s). Filters require the source and destination address to match a condition. Use optional keywords to configure finer granularity.
2. Apply the access filter to specified interfaces.

To create an IP-ACL, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# ip access-list List1 permit ip any any | Configures an IP-ACL called List1 and permits IP traffic from any source address to any destination address. |
| | switch(config)# no ip access-list List1 permit ip any any | Removes the IP-ACL called List1. |
| Step 3 | switch(config)# ip access-list List1 deny tcp any any | Updates List1 to deny TCP traffic from any source address to any destination address. |

To define an IP-ACL that restricts management access, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# ip access-list restrict_mgmt permit ip 10.67.16.0 0.0.0.255 any | Defines an entry in IP-ACL named restrict_mgmt allowing all addresses in the 10.67.16.0/24 subnet. |
| Step 3 | switch(config)# ip access-list restrict_mgmt permit icmp any any eq 8 | Adds an entry to IP-ACL named restrict_mgmt to allow any device to ping the MDS (icmp type 8). |
| Step 4 | switch(config)# ip access-list restrict_mgmt deny ip any any | Explicitly blocks all other access for access-list named restrict_mgmt. |

To use the operand and port options, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# ip access-list List2 deny tcp 1.2.3.0 0.0.0.255 eq port 5 any | Denies TCP traffic from 1.2.3.0 through source port 5 to any destination. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Adding filters to an Existing IP-ACL

After you create an IP-ACL, you place subsequent additions at the end of the IP-ACL. You cannot insert filters in the middle of an IP-ACL. Each configured entry is automatically added to the end of a IP-ACL.

To add entries to an existing IP-ACL, follow these steps:

| | Command | Purpose |
|--------|---|--------------------------------|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# ip access-list List1 permit tcp 10.1.1.2 0.0.0.0 172.16.1.1 0.0.0.0 eq port telnet | Permits TCP for Telnet traffic |
| | switch(config)# ip access-list List1 permit tcp 10.1.1.2 0.0.0.0 172.16.1.1 0.0.0.0 eq port http | Permits TCP for HTTP traffic. |
| | switch(config)# ip access-list List1 permit udp 10.1.1.2 0.0.0.0 172.16.1.1 0.0.0.0 | Permits UDP for all traffic. |
| | | |

Removing Entries from an Existing IP-ACL

To remove configured entries from an IP-ACL, follow these steps:

| | Command | Purpose |
|--------|---|-------------------------------------|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# no ip access-list List2 deny tcp 1.2.3.0 0.0.0.255 eq port 5 any | Removes this entry from the IP-ACL. |
| | switch(config)# no ip access-list x3 deny ip any any | Removes this entry from the IP-ACL. |
| | switch(config)# no ip access-list x3 permit ip any any | Removes this entry from the IP-ACL. |

Reading the IP-ACL Log Dump

Use the **log-deny** option at the end of an filter condition to log information about packets that match dropped entries. The log output displays the ACL number, permit or deny status, and port information.

For the input ACL, the log displays the raw MAC information. The keyword “MAC=” does not refer to showing an Ethernet MAC frame with MAC address information. It refers to the Layer 2 MAC-layer information dumped to the log. For the output ACL, the raw Layer 2 information is not logged.

The following is example of an input ACL log dump.

```
Jul 17 20:38:44 excal-2
%KERN-7-SYSTEM_MSG:
%IPACL-7-DENY:IN=vsan1 OUT=
MAC=10:00:00:05:30:00:47:df:10:00:00:05:30:00:8a:1f:aa:aa:03:00:00:00:08:00:45:00:00:54:00
:00:40:00:40:01:0e:86:0b:0b:0b:0c:0b:0b:0b:02:08:00:ff:9c:01:15:05:00:6f:09:17:3f:80:02:01
:00:08:09:0a:0b:0c:0d:0e:0f:10:11:12:13:14:15:16:17:18:19:1a:1b:1c:1d:1e:1f:20:21:22:23:24
:25:26:27:28:29:2a:2b SRC=11.11.11.12 DST=11.11.11.2 LEN=84 TOS=0x00 PREC=0x00 TTL=64 ID=0
DF PROTO=ICMP TYPE=8 CODE=0 ID=277 SEQ=1280
```

The following example is an output ACL log dump.

```
Jul 17 20:38:44 excal-2
%KERN-7-SYSTEM_MSG:
%IPACL-7-DENY:IN= OUT=vsan1 SRC=11.11.11.2 DST=11.11.11.2 LEN=84 TOS=0x00 PREC=0x00
TTL=255 ID=38095 PROTO=ICMP TYPE=0 CODE=0 ID=277 SEQ=1280
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Applying an IP-ACL to an Interface

You can define IP-ACLs without applying them. However, the IP-ACLs will have no effect until they are applied to an interface on the switch.




Tip

Apply the IP-ACL on the interface closest to the source of the traffic.

When you are trying to block traffic from source to destination, you can apply an inbound IP-ACL to M0 on Switch 1 instead of an outbound filter to M1 on Switch 3 (see [Figure 26-3](#)).

Figure 26-3 Denying Traffic on the Inbound Interface



The **access-group** option controls access to an interface. Each interface can only be associated with one IP-ACL per direction. The ingress direction can have a different IP-ACL than the egress direction. The IP-ACL becomes active on when applied to the interface.



Tip

Create all conditions in an IP-ACL before applying it to the interface.



Caution

If you apply an IP-ACL to an interface before creating it, all packets in that interface are dropped because the IP-ACL is empty.

The terms *in*, *out*, *source*, and *destination* are used as referenced by the switch.

- In—Traffic that arrives at the interface and which will go through the switch; the source is where it transmitted from and the destination is where it is transmitted to (on the other side of the router).



Tip

The IP-ACL applied to the interface for the ingress traffic affects both local and remote traffic.

- Out—Traffic that has already been through the switch and is leaving the interface;the source is where it transmitted from and the destination is where it is transmitted to.



Tip

The IP-ACL applied to the interface for the egress traffic only affects local traffic.

To apply an IP-ACL to an interface, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface mgmt0 switch(config-if)# | Configures a management interface (mgmt0). |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|---|--|
| Step 3 | switch(config-if)# ip access-group restrict_mgmt | Applies an IP-ACL called restrict_mgmt for both the ingress and egress traffic (default). |
| | switch(config-if)# no ip access-group NotRequired | Removes the IP-ACL called NotRequired. |
| Step 4 | switch(config-if)# ip access-group restrict_mgmt in | Applies an IP-ACL called restrict_mgmt (if it does not already exist) for ingress traffic. |
| | switch(config-if)# no ip access-group restrict_mgmt in | Removes the IP-ACL called restrict_mgmt for ingress traffic. |
| | switch(config-if)# ip access-group SampleName2 out | Applies an IP-ACL called SampleName (if it does not already exist) for local egress traffic. |
| | switch(config-if)# no ip access-group SampleName2 out | Remove the IP-ACL called SampleName for local egress traffic. |

IP-ACL Configuration Verification

Use the **show ip access-list** command to view the contents of configured access filters. Each access filter can have several conditions.

Example 26-1 Displays Configured IP-ACLs

```
switch# show ip access-list usage
Access List Name/Number      Filters IF   Status      Creation Time
-----
abc                          3          7    active    Tue Jun 24 17:51:40 2003
x1                           3          1    active    Tue Jun 24 18:32:25 2003
x3                           0          1  not-ready Tue Jun 24 18:32:28 2003
```

Example 26-2 Displays a Summary of the Specified IP-ACL

```
switch# show ip access-list abc
ip access-list abc permit tcp any any (0 matches)
ip access-list abc permit udp any any (0 matches)
ip access-list abc permit icmp any any (0 matches)
ip access-list abc permit ip 10.1.1.0 0.0.0.255 (2 matches)
ip access-list abc permit ip 10.3.70.0 0.0.0.255 (7 matches)
```

IP-ACL Counter Cleanup

Use the **clear** command to clear the counters for a specified IP-ACL entry.



Note

You cannot use this command to clear the counters for each individual filter.

```
switch# clear ip access-list counters abc
```

Send documentation comments to mdsfeedback-doc@cisco.com.

IPFC Configuration

Once the VSAN interface is created, you can specify the IP address for that VSAN.

Configuring an IP Address in a VSAN

To configure a VSAN interface and an IP address for that interface, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface vsan 1 switch(config-if)# | Configures the interface for the specified VSAN (1). |
| Step 3 | switch(config-if)# ip address 10.0.0.12 255.255.255.0 switch(config-if)# | Configures the IP address and netmask for the selected interface. |

Enabling IP Routing

By default, the IP routing feature is disabled in all switches.

To enable the IP routing feature, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# ip routing switch(config)# | Enables IP routing (disabled by default). |
| Step 3 | switch(config)# no ip routing switch(config)# | Disables IP routing and reverts to the factory settings. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Configuring IP Static Routes

Static routing is a mechanism to configure IP routes on the switch. You can configure more than one static route.

If your configuration does not need an external router, you can use static routing.

If a VSAN has multiple exit points, configure static routes to direct traffic to the appropriate gateway switch. IP routing is disabled by default on any gateway switch between the out-of-band management interface and the default VSAN, or between directly connected VSANs.

To configure a static route, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# confi g t | Enters configuration mode. |
| Step 2 | switch(config)# IP route <network IP address> <netmask> <next hop IP address> <distance> <number> interface <vsan number> For example: switch(config)# IP route 10.0.0.0 255.0.0.0 20.20.20.10 distance 10 interface vsan 1 switch(config)# | Configures the static route for the specified IP address, subnet mask, next hop, and distance, and VSAN or management interface. |

Displaying and Clearing ARPs

Address Resolution Protocol (ARP) entries in Cisco MDS 9000 Family switches can be displayed, deleted, or cleared. The ARP feature is enabled on all switches.

- Use the **show arp** command to display the ARP table.

```
switch# show arp
Protocol Address          Age (min)    Hardware Addr  Type  Interface
Internet  171.1.1.1              0            0006.5bec.699c  ARPA  mgmt0
Internet  172.2.0.1              4            0000.0c07.ac01  ARPA  mgmt0
```

- Use the **no arp** command in configuration mode to remove an ARP entry from the ARP table.

```
switch(config)# no arp 172.2.0.1
```

- Use the **clear arp** command to delete all entries from the ARP table. The ARP table is empty by default.

```
switch# clear arp-cache
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying IP Interface Information

Use the following **show** commands to view configured IP interface information (see Examples 26-3 to 26-6).

Example 26-3 *Displays the VSAN Interface*

```
switch# show interface vsan1
vsan1 is up, line protocol is up
  WWPN is 10:00:00:05:30:00:59:1f, FCID is 0x9c0100
  Internet address is 10.1.1.1/24
  MTU 1500 bytes, BW 1000000 Kbit
  0 packets input, 0 bytes, 0 errors, 0 multicast
  0 packets output, 0 bytes, 0 errors, 0 dropped
```



Note You can see the output for this command only if you have previously configured a virtual network interface (see the [“Configuring an IP Address in a VSAN”](#) section on page 26-12).

Example 26-4 *Displays the Connected and Static Route Details*

```
switch# show ip route

Codes: C - connected, S - static

Default gateway is 172.22.95.1

C 172.22.95.0/24 is directly connected, mgmt0
C 10.1.1.0/24 is directly connected, vsan1
```

Example 26-5 *Displays Configured Routes*

```
switch# show ip route configured
```

| Destination | Gateway | Mask | Metric | Interface |
|-------------|-------------|---------------|--------|-----------|
| default | 172.22.95.1 | 0.0.0.0 | 0 | mgmt0 |
| 10.1.1.0 | 0.0.0.0 | 255.255.255.0 | 0 | vsan1 |
| 172.22.95.0 | 0.0.0.0 | 255.255.255.0 | 0 | mgmt0 |

Example 26-6 *Displays the IP Routing Status*

```
switch# show ip routing
ip routing is disabled
```

Send documentation comments to mdsfeedback-doc@cisco.com.

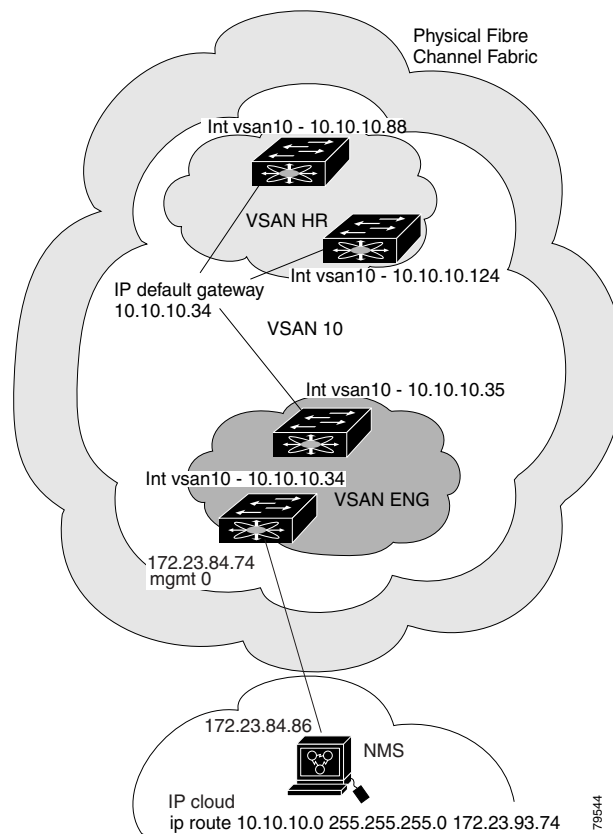
Overlay VSAN Configuration

VSANs enable deployment of larger SANs by overlaying multiple logical SANs, each running its own instance of fabric services, on a single large physical network. This partitioning of fabric services reduces network instability by containing fabric reconfiguration and error conditions within an individual VSAN. VSANs also provide the same isolation between individual VSANs as physically separated SANs. Traffic cannot cross VSAN boundaries and devices may not reside in more than one VSAN. Because each VSAN runs separate instances of fabric services, each VSAN has its own zone server and can be zoned in exactly the same way as SANs without VSAN capability.

To configure an overlay VSAN, follow these steps:

- Step 1** Add the VSAN to the VSAN database on all switch in the fabric.
- Step 2** Create a VSAN interface for the VSAN on all switches in the fabric. Any VSAN interface belonging to the VSAN has an IP address in the same subnet. Create a route to the IPFC cloud on the IP side
- Step 3** Configure a default route on every switch in the Fibre Channel fabric pointing to the switch that provides NMS access.
- Step 4** Configure default gateway (route) and the IP address on switches that point to the NMS (see [Figure 26-4](#)).

Figure 26-4 Overlay VSAN Configuration Example



Send documentation comments to mdsfeedback-doc@cisco.com.

**Note**

To configure the management interface displayed in [Figure 26-4](#), set the default gateway to an IP address on the Ethernet network.

The following procedure configures an overlay VSAN in one switch. This procedure must be repeated for each switch in the fabric.

To configure an overlay VSAN in one switch (using the example in [Figure 26-4](#)), follow these steps:

| | Command | Purpose |
|---------------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# vsan database switch-config-vsan-db# | Configures the VSAN database. |
| Step 3 | switch--config-vsan-db# vsan 10 name MGMT_VSAN | Defines the VSAN in the VSAN database on all of the switches in the Fibre Channel fabric. |
| Step 4 | switch--config-vsan-db# exit switch(config)# | Exits the VSAN database mode. |
| Step 5 | switch(config)# interface vsan10 switch(config-if)# | Creates a VSAN interface (VSAN 10). |
| Step 6 | switch(config-if)# ip address 10.10.10.0 netmask 255.255.255.0 | Assigns an IP address and netmask for this switch. |
| Step 7 | switch(config-if)# no shut | Enables the configured interface. |
| Step 8 | switch--config-if# end switch# | Exits to EXEC mode. |
| Step 9 | switch# exit | Exits the switch and returns to the NMS. In this example the NMS is assumed to be on the same subnet of the Ethernet management interface of the edge that provides access to the Fibre Channel fabric. |

To configure the NMS station displayed in [Figure 26-4](#), follow this step:

| | Command | Purpose |
|---------------|--|---|
| Step 1 | nms# route ADD 10.10.10.0 MASK 255.255.255.0 172.22.93.74 | Defines a static route on the NMS pointing to the management interface of the edge switch that provides access to the Fibre Channel fabric. |

Send documentation comments to mdsfeedback-doc@cisco.com.

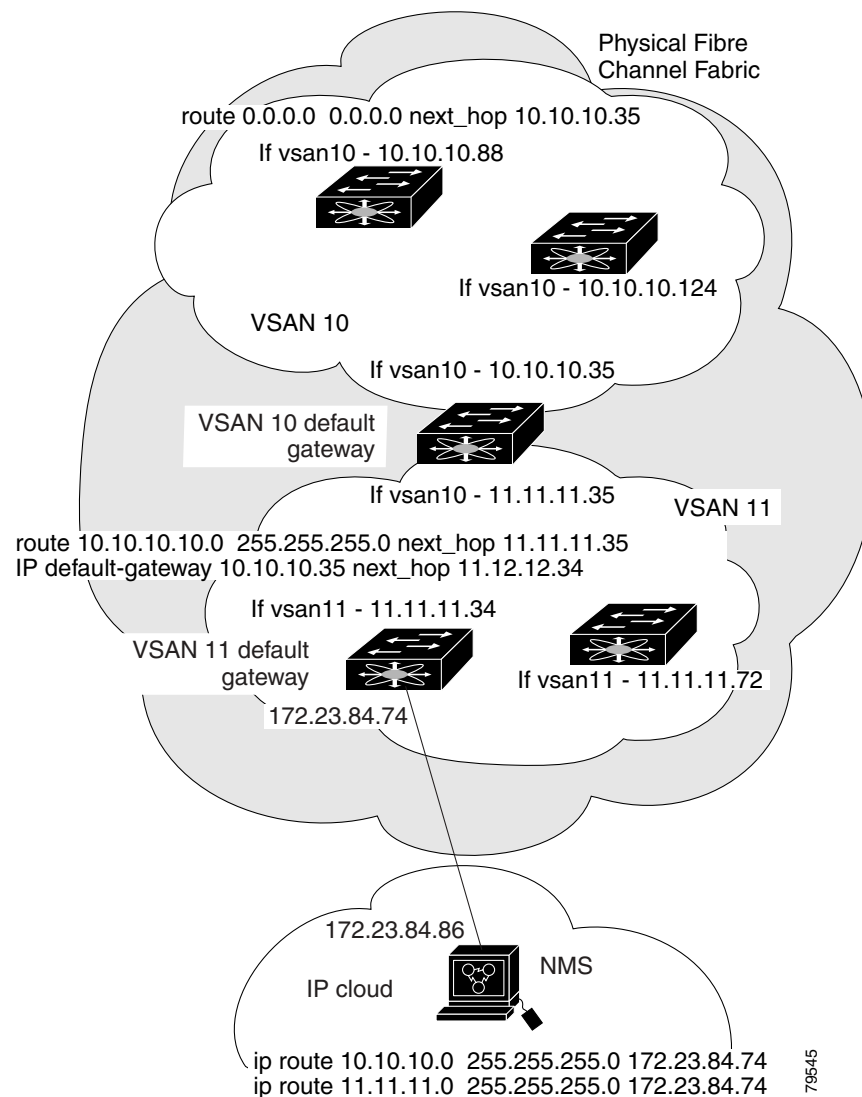
Multiple VSAN Configuration

More than one VSAN can be used to segment the management network in multiple subnets. An active interface must be present on the switch for the VSAN interface to be enabled.

To configure multiple VSANs, follow these steps:

- Step 1** Add the VSAN to the VSAN database on any switch in the fabric.
- Step 2** Create a VSAN interface for the appropriate VSAN on any switch in the fabric.
- Step 3** Assign an IP address on every VSAN interface on the same subnet as the corresponding VSAN.
- Step 4** Define the multiple static route on the Fibre Channel switches and the IP cloud (see [Figure 26-5](#)).

Figure 26-5 Multiple VSANs Configuration Example



Send documentation comments to mdsfeedback-doc@cisco.com.

To configure an overlay VSAN (using the example in [Figure 26-5](#)), follow these steps:

| | Command | Purpose |
|---------|---|---|
| Step 1 | switch# confi g t | Enters configuration mode. |
| Step 2 | switch(config)# vsan database switch-config-vsan-db# | Configures the VSAN database. |
| Step 3 | switch-config-vsan-db# vsan 10 name MGMT_VSAN_10 switch-config-vsan-db# | Defines the VSAN in the VSAN database on all of the switches in VSAN 10. |
| Step 4 | switch-config-vsan-db# exit switch(config)# | Exits the database 10 mode. |
| Step 5 | switch-config-vsan-db# vsan 11 name MGMT_VSAN_11 switch-config-vsan-db# | Defines the VSAN in the VSAN database on all of the switches in VSAN 11. |
| Step 6 | switch-config-vsan-db# exit switch(config)# | Exits the VSAN database 11 mode. |
| Step 7 | switch(config)# interface vsan10 switch(config-if)# | Enters the VSAN 10 interface configuration mode for VSAN 10. |
| Step 8 | switch(config-if)# ip address 10.10.10.0 netmask 255.255.255.0 switch(config-if)# | Assigns an IP address and netmask for this switch. |
| Step 9 | switch(config-if)# no shut | Enables the configured interface for VSAN 10. |
| Step 10 | switch--config-if# exit switch(config)# | Exits the VSAN 10 interface mode. |
| Step 11 | switch(config)# interface vsan11 switch(config-if)# | Enters the VSAN 11 interface configuration mode. |
| Step 12 | switch(config-if)# ip address 11.11.11.0 netmask 255.255.255.0 switch(config-if)# | Assigns an IP address and netmask for this switch |
| Step 13 | switch(config-if)# no shut | Enables the configured interface for VSAN 11. |
| Step 14 | switch-config-if# end switch# | Exits to EXEC mode. |
| Step 15 | switch# exit | Exits the switch and returns to the NMS. In this example the NMS is assumed to be on the same subnet of the Ethernet management interface of the edge that provides access to the Fibre Channel fabric. |
| Step 16 | NMS# route ADD 10.10.10.0 MASK 255.255.255.0 172.22.93.74 | Defines a static route on the NMS pointing to the management interface of the edge switch that provides access to the IP cloud. |
| Step 17 | NMS# route ADD 11.11.11.0 MASK 255.255.255.0 172.22.93.74 | Defines a static route for VSAN 11 on the NMS pointing to the management interface of the edge switch that provides access to the Fibre Channel fabric. |
| Step 18 | switch# route 10.10.10.0 255.255.255.0 next_hop 11.11.11.35 | Defines the route to reach subnet 10 from subnet 11. |

Send documentation comments to mdsfeedback-doc@cisco.com.

The Virtual Router Redundancy Protocol

Cisco MDS 9000 Family switches are compliant with RFC 2338 standards for Virtual Router Redundancy Protocol (VRRP) features. This section provides details on the VRRP feature.

VRRP Features


VRRP provides a redundant alternative path to the gateway switch, which has connectivity to the NMS. VRRP has the following characteristics and advantages:

- VRRP is a restartable application.
- When a VRRP master fails, the VRRP backup takes over within three times the advertisement time.
- VRRP over Ethernet, VRRP over VSAN, and Fibre Channel functions are implemented as defined in RFC 2338.
- A virtual router is mapped to each VSAN and Ethernet interface with its unique virtual router IP, virtual router MAC, and VR ID.
- VR IDs can be reused in multiple VSANs with a different virtual router IP mapping.
- Up to 255 virtual router groups can be assigned in each VSAN.
- VRRP security provides three options, including no authentication, simple text authentication, and MD5 authentication.

VRRP Functionality

In [Figure 26-6](#), switch A is the VRRP master and switch B is the VRRP backup switch. Both switches have IP address to VRRP mapping configured. The other switches set switch A as the default gateway. If switch A fails, the other switches don't have to change the routing configurations as switch B automatically becomes the master and takes over the function of a gateway.

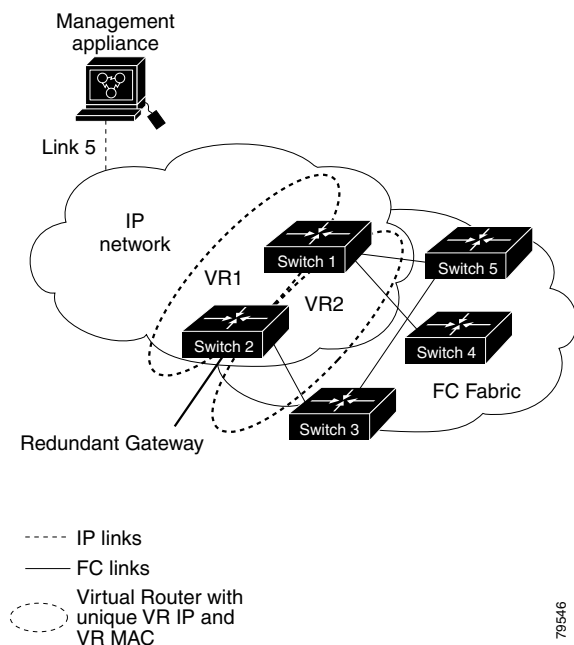
Figure 26-6 VRRP Functionality



Send documentation comments to mdsfeedback-doc@cisco.com.

In Figure 26-7, the fabric example has two virtual router groups (VR1 and VR 2) because a virtual router cannot span across different types of interfaces. In both switch 1 and switch 2, the Ethernet interface is in VR 1 and the FC interface is in VR 2. Each virtual router is uniquely identified by the VSAN interface and the VR ID.

Figure 26-7 Redundant Gateway



Virtual Router Addition and Deletion

All VRRP configurations should be replicated across switches in a fabric that runs VRRP.

To create or remove a VR, follow these steps:

| | Command | Purpose |
|--------|---|---------------------------------------|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# interface vsan 1 switch(config-if)# | Configures a VSAN interface (VSAN 1). |
| Step 3 | switch(config-if)# vrrp 250 switch(config-if-vrrp) | Creates a VR ID 250. |
| | switch(config-if-vrrp)# no vrrp 250 switch(config-if) | Removes a VR ID 250. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Virtual Router Initiation

By default, a virtual router is always disabled. VRRP can be configured only if this state is enabled. Be sure to configure at least one IP address before attempting to enable a VR.

To enable or disable a virtual router, follow these steps:

| | Command | Purpose |
|--------|--|------------------------------|
| Step 1 | <code>switch(config-if-vrrp)# no shutdown</code> | Enables VRRP configuration. |
| | <code>switch(config-if-vrrp)# shutdown</code> | Disables VRRP configuration. |

Virtual Router IP Address Addition

One primary IP address and multiple secondary addresses can be configured for a switch. If the configured IP address is the same as the interface IP address, this switch automatically owns the IP address.

To configure an IP address for a virtual router, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | <code>switch# config t</code> | Enters configuration mode. |
| Step 2 | <code>switch(config)# interface vsan 1</code> <code>switch(config-if)#</code> | Configures a VSAN interface (VSAN 1). |
| Step 3 | <code>switch(config-if)# interface ip address 10.0.0.12 255.255.255.0xi</code> | Configures an IP address. The IP address must be configured before the VRRP is added. |
| Step 4 | <code>switch(config-if)# vrrp 250</code> <code>switch(config-if-vrrp)#</code> | Creates VR ID 250. |
| Step 5 | <code>switch(config-if-vrrp)# address 10.0.0.10</code> | Configures the IP address (10.0.0.10) for the selected VR. Note This IP address should be in the same subnet as the IP address of the interface. |
| | <code>switch(config-if-vrrp)# no address 10.0.0.10</code> | Removes the IP address (10.0.0.10) for the selected VR. |

Priority for the Virtual Router

The valid range to assign a virtual router priority is 1 to 254 with 1 being the lowest priority and 254 being the highest priority. The default value is 100 for switches with secondary IP addresses and 255 for a switch with the primary IP address.

To set the priority for a virtual router, follow these steps:

| | Command | Purpose |
|--------|--|---------------------------------------|
| Step 1 | <code>switch# config t</code> | Enters configuration mode. |
| Step 2 | <code>switch(config)# interface vsan 1</code> <code>switch(config-if)#</code> | Configures a VSAN interface (VSAN 1). |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|--|--|
| Step 3 | switch(config-if)# vrrp 250 switch(config-if-vrrp)# | Creates a virtual router. |
| Step 4 | switch(config-if-vrrp)# priority 2 switch(config-if-vrrp)# | Configures the priority for the selected VRRP. |
| | | Note Priority 255 cannot be preempted. |

Time Interval for Advertisement Packets

The valid time range for an advertisement packet is between 1 and 255 seconds with the default being 1 (one) second. If the switch has the primary IP address, this time must be specified.

To set the priority for a virtual router, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface vsan 1 switch(config-if)# | Configures a VSAN interface (VSAN 1). |
| Step 3 | switch(config-if)# vrrp 250 switch(config-if-vrrp)# | Creates a virtual router. |
| Step 4 | switch(config-if-vrrp)# advertisement-interval 15 | Sets the interval time in seconds between sending advertisement frames. |



Note

If the virtual IP address is also the IP address for the interface, then preemption is implicitly applied.



Note

The VRRP **preempt** option is not supported on IP storage Gigabit Ethernet interfaces.

To enable or disable preempting, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface vsan 1 switch(config-if)# | Configures a VSAN interface (VSAN 1). |
| Step 3 | switch(config-if)# vrrp 250 switch(config-if-vrrp)# | Creates a virtual router. |
| Step 4 | switch(config-if-vrrp)# preempt | Enables the higher priority backup virtual router to preempt the lower priority master virtual router. |
| | | Note This preemption does not apply to the primary IP address. |
| | switch(config-if-vrrp)# no preempt | Disables the preempt option and allows the master to keep its priority level. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Virtual Router Authentication

VRRP security provides three options, including simple text authentication, MD5 authentication, and no authentication.

- Simple text authentication uses a unique, 1 to 8 character password that is used by all switches participating in the same virtual router. This password should be different from other security passwords.
- MD5 authentication uses a unique, 16 character key that is shared by all switches participating in the same virtual router. This secret key is shared by all switches in the same virtual router.
- No authentication is the default option.

You can configure the key using the authentication option in the VRRP submode and distribute it using the configuration file. The security parameter index (SPI) settings assigned in this option should be unique for each VSAN.



Note

All VRRP configurations must be duplicated.

To set an authentication option for a virtual router, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface vsan 1 switch(config-if)# | Configures a VSAN interface (VSAN 1). |
| Step 3 | switch(config-if)# vrrp 250 switch(config-if-vrrp)# | Creates a virtual router. |
| Step 4 | switch(config-if-vrrp)# authentication text password | Assigns the simple text authentication option and specifies the password for this option. |
| | switch(config-if-vrrp)# authentication md5 password2003 spi 0x2003 | Assigns MD5 authentication option and specifies the key and the unique SPI value for this option. The SPI and the valid range is 0x100 to 0xFFFFFFFF. |
| | switch(config-if-vrrp)# no authentication | Assigns the no authentication option, which is the default. |

Priority Based on Interface State

The tracking feature is disabled by default. When you specify the tracking option, the priority of the virtual router is changed based on the state of another interface in the switch. When the tracked interface is down, the priority of the virtual router is changed to a lower priority value. When the tracked interface is up, the priority of the virtual router is restored to its original value. You can track one of two interfaces on a switch in the Cisco MDS 9000 Family: a specified VSAN interface or a management interface.

To track the interface priority for a virtual router, follow these steps:

| | Command | Purpose |
|--------|---|---------------------------------------|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface vsan 1 switch(config-if)# | Configures a VSAN interface (VSAN 1). |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|--|---|
| Step 3 | switch(config-if)# vrrp 250 switch(config-if-vrrp)# | Creates a virtual router. |
| Step 4 | switch(config-if-vrrp)# track interface mgmt 0 priority 2 | Specifies the priority of the virtual router to be modified based on the state of the management interface. |
| | switch(config-if-vrrp)# no track | Disables the tracking feature. |

Displaying VRRP Information

Use the **show vrrp vr** command to display configured VRRP information (see Examples 26-7 to 26-10).

Example 26-7 Displays VRRP Configured Information

```
switch# show vrrp vr 7 interface vsan 2 configuration
vr id 7 configuration
admin state down
priority 100
no authentication
advertisement-Interval 1
preempt yes
tracking interface vsan1 priority 2
protocol IP
```

Example 26-8 Displays VRRP Status Information

```
switch# show vrrp vr 7 interface vsan 2 status
vr id 7 status
MAC address 00:00:5e:00:01:07
Operational state: init
```

Example 26-9 Displays VRRP Statistics

```
switch# show vrrp vr 7 interface vsan 2 statistics
vr id 7 statistics
Become master 0
Advertisement 0
Advertisement Interval Error 0
Authentication Failure 0
TTL Error 0
Priority 0 Received 0
Priority 0 Sent 0
Invalid Type 0
Mismatch Address List 0
Invalid Authentication Type 0
Mismatch Authentication 0
Invalid Packet Length 0
```

Example 26-10 Displays VRRP Cumulative Statistics

```
switch# show vrrp statistics
Invalid checksum 0
Invalid version 0
Invalid VR ID 0
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Clearing VRRP Statistics

Use the **clear vrrp** command to clear all the software counters for the specified virtual router (see [Example 26-11](#)).

Example 26-11 Clears VRRP Information

```
switch# clear vrrp 7 interface vsan2
switch#
```

DNS Server Configuration

The DNS client on the switch communicates with the DNS server to perform the IP address-name server correspondence.

The DNS server may be dropped after two attempts due to one of the following reasons:

- The IP address or the switch name is wrongly configured
- The DNS server is not reachable due to external reasons (reasons beyond our control)



Note

When accessing a Telnet host, if the DNS server is not reachable (for any reason) the switch login prompt may take a longer time to appear. If so, verify that the DNS server is accurately configured and reachable.

To configure a DNS server, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ip domain-lookup | Enables the IP Domain Naming System (DNS)-based host name-to-address translation. |
| | switch(config)# no ip domain-lookup | Disables (default) the IP DNS-based host name-address translation and reverts to the factory default. |
| Step 3 | switch(config)# no ip domain-name cisco.com | Disables the domain name and reverts to the factory default. |
| | switch(config)# ip domain-name cisco.com | Enables (default) the default domain name feature used to complete unqualified host names. Any IP host name that does not contain a domain name (that is, any name without a dot), will have the dot and cisco.com appended to it before being added to the host table. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|--|---|
| Step 4 | switch(config)# ip domain-list harvard.edu | Defines a filter of default domain names to complete unqualified host names, use the ip domain-list global configuration command. You can define up to 10 domain names in this filter. To delete a name from a filter, use the no form of this command. |
| | switch(config)# ip domain-list stanford.edu | |
| | switch(config)# ip domain-list yale.edu | |
| | switch(config)# no ip domain-list | Deletes the defined filter and reverts to factory default. No domains are configured by default. |
| | Note If you have not configured a domain list, the domain name that you specified with the ip domain-name global configuration command is used. If you did configure a domain list, the default domain name is not used. The ip domain-list command is similar to the ip domain-name command, except that with the ip domain-list command you can define a list of domains, each to be tried in turn. | |
| Step 5 | switch(config)# ip name-server 15.1.0.1 15.2.0.0 | Specifies the first address (15.1.0.1) as the primary server and the second address (15.2.0.0) as the secondary sever. You can configure a maximum of six servers. |
| | switch(config)# no ip name-server | Deletes the configured server(s) and reverts to factory default. No server is configured by default. |
| | Note Alternatively, you can configure the DNS entry using the switch names (instead of IP addresses). The configured switch name automatically looks up the corresponding IP address. | |

Displaying DNS Host Information

Use the **show hosts** command to display the DNS configuration (see [Example 26-12](#)).

Example 26-12 Displays Configured Host Details

```
switch# show hosts
Default domain is cisco.com
Domain list: ucsc.edu harvard.edu yale.edu stanford.edu
Name/address lookup uses domain service
Name servers are 15.1.0.1 15.2.0.0
```

Default Settings

[Table 26-3](#) lists the default settings for FSPF features.

Table 26-3 Default FSPF Settings

| Parameters | Default |
|--|--------------------------------------|
| FSPF | Enabled on all E ports and TE ports. |
| SPF computation | Dynamic. |
| SPF hold time | 0. |
| Backbone region | 0. |
| Acknowledgment interval (RxmtInterval) | 5 seconds. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 26-3 **Default FSPF Settings (continued)**

| Parameters | Default |
|-------------------------------|---|
| Refresh time (LSRefreshTime) | 30 minutes. |
| Maximum age (MaxAge) | 60 minutes. |
| Hello interval | 20 seconds. |
| Dead interval | 80 seconds. |
| Distribution tree information | Derived from the principal switch (root node). |
| Routing table | FSPF stores up to 16 equal cost paths to a given destination. |
| Load balancing | Based on destination ID and source ID on different, equal cost paths. |
| In-order delivery | Disabled. |
| Drop latency | Disabled. |
| Static route cost | If the cost (metric) of the route is not specified, the default is 10. |
| Remote destination switch | If the remote destination switch is not specified, the default is direct. |
| Multicast routing | Uses the principal switch to compute the multicast tree. |

Send documentation comments to mdsfeedback-doc@cisco.com.



Configuring FICON

Fibre Connection (FICON) interface capabilities enhance the Cisco MDS 9000 Family by supporting both open systems and mainframe storage network environments. Inclusion of Control Unit Port (CUP) support further enhances the MDS offering by allowing in-band management of the switch from FICON processors.

The fabric binding feature helps prevent unauthorized switches from joining the fabric or disrupting current fabric operations. The Registered Link Incident Report ((RLIR) application provides a method for a switchport to send a LIR to a registered Nx-port.

This chapter includes the following sections:


- [About FICON, page 27-2](#)
- [MDS-Specific FICON Advantages, page 27-3](#)
- [FICON Port Numbering, page 27-7](#)
- [FICON VSAN Prerequisites, page 27-11](#)
- [Enabling FICON, page 27-11](#)
- [You can enable FICON on a per VSAN basis in one of three ways:, page 27-11](#)
- [Manually Enabling FICON, page 27-15](#)
- [Running Configuration Automatic Save, page 27-19](#)
- [Binding Port Numbers to PortChannels, page 27-20](#)
- [Binding Port Numbers to FCIP Interfaces, page 27-20](#)
- [Configuring FICON Ports, page 27-21](#)
- [FICON Configuration Files, page 27-23](#)
- [Port Swapping, page 27-25](#)
- [Moving a FICON VSAN to an Offline State, page 27-27](#)
- [Clearing FICON Device Allegiance, page 27-27](#)
- [CUP In-Band Management, page 27-27](#)
- [Displaying FICON Information, page 27-28](#)
- [Fabric Binding Configuration, page 27-37](#)
- [Displaying RLIR Information, page 27-44](#)
- [Default Settings, page 27-48](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

About FICON

The Cisco MDS 9000 Family supports the Fibre Channel Protocol (FCP), FICON, iSCSI, and FCIP capabilities within a single, high availability platform. This solution simplifies purchasing, reduces deployment and management costs, and reduces the complex evolution to shared mainframe and open systems storage networks (see [Figure 27-1](#)).

Figure 27-1 Shared System Storage Network



FCP and FICON are different FC4 protocols and their traffic are independent of each other. If required, devices using these protocols can be isolated using VSANs.

FICON Requirements

The FICON feature has the following requirements:

- FICON features can be implemented in the following switches running Cisco MDS SAN-OS Release 1.3(4a) or later:
 - Any switch in the Cisco MDS 9500 Series.
 - Any switch in the Cisco MDS 9200 Series.



Note

The FICON feature is not supported on Cisco MDS 9120 and 9140 switches or the 32-port Fibre Channel switching module.

- You need the MAINFRAME_PKG license to configure FICON parameters (see [Chapter 3, “Obtaining and Installing Licenses”](#)).

Send documentation comments to mdsfeedback-doc@cisco.com.

MDS-Specific FICON Advantages

This section explains the additional FICON advantages in Cisco MDS switches.

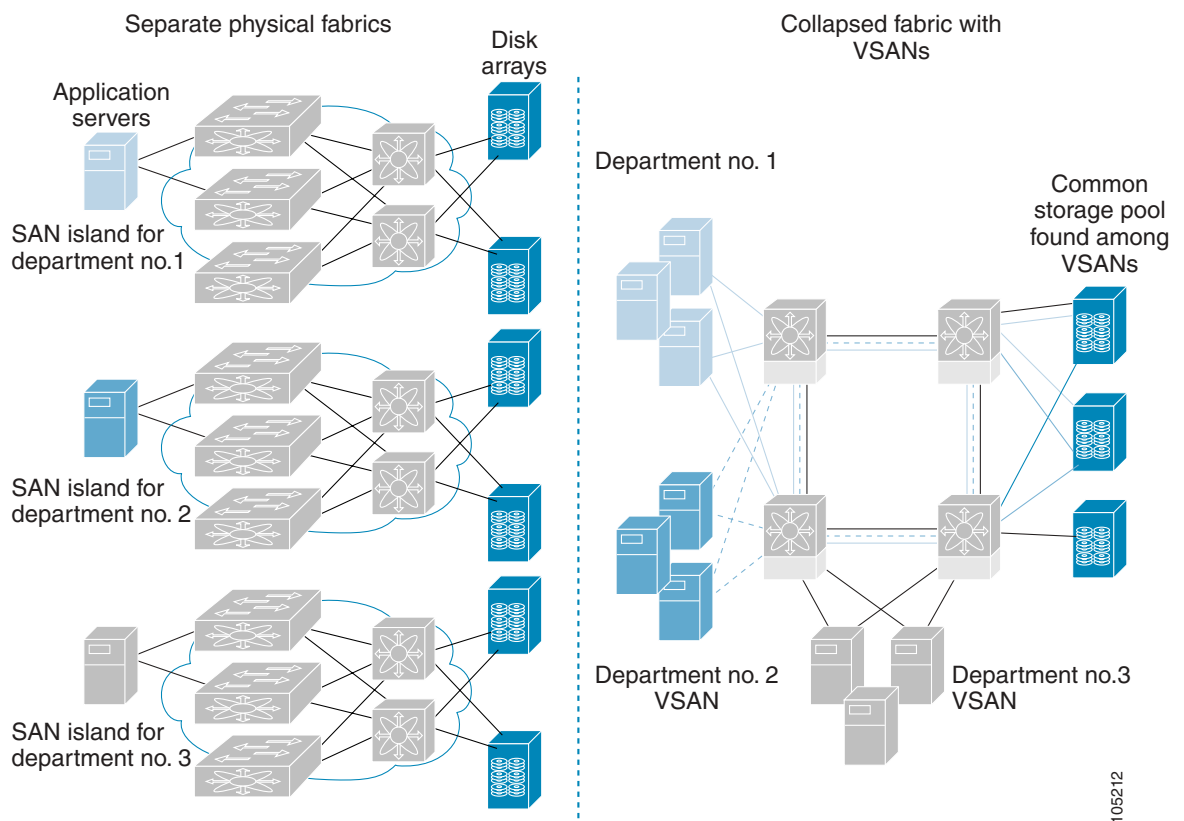
Fabric Optimization with VSANs

Generally, separate physical fabrics have a high level of switch management and have a higher implementation cost. Further, the ports in each island may be over-provisioned depending on the fabric configuration.

By using the Cisco MDS-specific VSAN technology, you can introduce greater efficiency between these physical fabrics by lowering the cost of over-provisioning and reducing the number of switches to be managed.

VSANs also help you to move unused ports nondisruptively and provide a common redundant physical infrastructure (see [Figure 27-2](#)).

Figure 27-2 VSAN-Specific Fabric Optimization



VSANs enable global SAN consolidation by allowing you to convert existing SAN islands into virtual SAN islands on a single physical network. It provides hardware-enforced security and separation between applications or departments to allow coexistence on a single network. It also allows virtual rewiring to consolidate your storage infrastructure. You can move assets between departments or applications without the expense and disruption of physical relocation of equipment.

Send documentation comments to mdsfeedback-doc@cisco.com.

**Note**

While you can configure up to 256 VSANs in any Cisco MDS switch, you can enable FICON in eight of these VSANs.

FCIP Support

The multilayer architecture of the Cisco MDS 9000 Family enables a consistent feature set over a protocol-agnostic switch fabric. Cisco MDS 9500 Series and 9200 Series switches transparently integrate Fibre Channel, FICON, and Fibre Channel over IP (FCIP) in one system. The FICON over FCIP feature enables cost-effective access to remotely located mainframe resources. With the Cisco MDS 9000 Family platform, storage replication services such as IBM PPRC and XRC can be extended over metro to global distances using ubiquitous IP infrastructure and simplifying business continuance strategies.

**Caution**

When write-acceleration is enabled in an FCIP interface, a FICON VSAN will not be enabled in that interface. Likewise, if a FCIP interface is up in a FICON VSAN, write-acceleration cannot be enabled on that interface.

See [Chapter 28, “Configuring IP Storage”](#) for more information on FCIP.

PortChannel Support

The Cisco MDS implementation of FICON provides support for efficient utilization and increased availability of inter-switch links necessary to build stable large-scale SAN environments. PortChannels ensure an enhanced ISL availability and performance in Cisco MDS switches.

See [Chapter 14, “Configuring PortChannels”](#) for more information on PortChannels.

VSANs for FICON and FCP Intermixing

Cisco MDS 9000 Family FICON-enabled switches simplify deployment of even the most complex intermix environments. Multiple logical FICON, Z-Series Linux/FCP, and Open-Systems FCP fabrics can be overlaid onto a single physical fabric by simply creating VSANs as required for each service. VSANs provide both hardware isolation and protocol specific fabric services, eliminating the complexity and potential instability of zone-based intermix schemes.

By default, the FICON feature is disabled in all switches in the Cisco MDS 9000 Family. When the FICON feature is disabled, FC IDs can be allocated seamlessly. Intermixed environments are addressed by the Cisco SAN-OS software. The challenge of mixing Fibre Channel Protocol (FCP) and FICON protocols are addressed by Cisco MDS switches when implementing VSANs.

Switches and directors in the Cisco MDS 9000 Family support FCP and FICON protocol intermixing at the port level. If these protocols are intermixed in the same switch, you can use VSANs to isolate FCP and FICON ports.

**Tip**

When creating an intermix environment, place all FICON devices in one VSAN (other than the default VSAN) and segregate the FCP switch ports in a separate VSAN (other than the default VSAN). This isolation ensures proper communication for all connected devices.

Send documentation comments to mdsfeedback-doc@cisco.com.

Cisco MDS-Supported FICON Features

The Cisco MDS 9000 Family FICON features include:

- Flexibility and investment protection—The Cisco MDS 9000 Family shares common switching and service modules across the Cisco MDS 9500 Series and the 9200 Series.

Refer to the *Cisco MDS 9500 Series Hardware Installation Guide* and the *Cisco MDS 9200 Series Hardware Installation Guide*).

- High-availability FICON-enabled director—The Cisco MDS 9500 Series combines nondisruptive software upgrades, stateful process restart and failover, and full redundancy of all major components for a new standard in director-class availability. It supports up to 224 autosensing, 2/1-Gbps, FICON or Fibre Channel FCP ports in any combination in a single chassis and up to 768 Fibre Channel ports in a single rack. The 1.44 Tbps of internal system bandwidth ensures smooth integration of future 10-Gbps modules. See [Chapter 5, “Configuring High Availability.”](#)
- Infrastructure protection—Common software releases infrastructure protection is available across all Cisco MDS 9000 platforms. See [Chapter 6, “Software Images.”](#)
- VSAN technology—The Cisco MDS 9000 Family introduces VSAN technology for hardware-enforced, isolated environments within a single physical fabric for secure sharing of physical infrastructure and enhanced FICON intermix support. See [Chapter 10, “Configuring and Managing VSANs.”](#)
- Port-level configurations—BB_credits, beacon mode, and port security for each port. See the [“Buffer-to-Buffer Credits” section on page 12-12](#), [“Identifying the Beacon LEDs” section on page 12-17](#), and [Chapter 13, “Configuring Trunking.”](#)
- Alias name configuration—instead of the WWN, for switches and attached node devices. See [Chapter 15, “Configuring and Managing Zones.”](#)
- Comprehensive security framework—The Cisco MDS 9000 Family supports RADIUS authentication, Simple Network Management Protocol Version 3 (SNMPv3), role-based access control, Secure Shell Protocol (SSH), Secure File Transfer Protocol (SFTP), VSANs, hardware-enforced zoning, ACLs, fabric binding, Fibre Channel Security Protocol (FC-SP), LUN zoning, read-only zones, and VSAN-based access control. See [Chapter 19, “Configuring Switch Security”](#) and [Chapter 20, “Configuring Fabric Security.”](#)
- Traffic encryption—IPSec is supported over FCIP. You can encrypt FICON and Fibre Channel traffic that is carried over FCIP. See [Chapter 29, “Configuring IPsec Network Security.”](#)
- View the local accounting log to locate FICON events. See the [“Local AAA Services” section on page 19-19](#).
- Unified storage management—Cisco MDS 9000 FICON-enabled switches are fully IBM CUP standard compliant for in-band management using the IBM S/A OS/390 I/O operations console. See the [“CUP In-Band Management” section on page 27-27](#).
- Port address-based configurations—port name, blocked or unblocked state, and the prohibit connectivity attributes. See the [“Configuring FICON Ports” section on page 27-21](#).
- Display the following information:
 - Individual Fibre Channel ports, such as the port name, port number, Fibre Channel address, operational state, type of port, and login data.
 - Nodes attached to ports.
 - Port performance and statistics.

See the [“Displaying FICON Information”](#) section in this chapter.

Send documentation comments to mdsfeedback-doc@cisco.com.


- Store and apply configuration files. See the [“FICON Configuration Files”](#) section on page 27-23.
- FICON and Open Systems Management Server features if installed. See the [“VSANs for FICON and FCP Intermixing”](#) section on page 27-4.
- Enhanced Cascading Support. See the [“CUP In-Band Management”](#) section on page 27-27.
- Set the date and time on the switch. See the [“FICON Host Control”](#) section on page 27-17.
- Configure SNMP trap recipients and community names. See the [“FICON SNMP Control”](#) section on page 27-18.
- Call Home configurations—director name, location, description, and contact person. See [Chapter 30, “Call Home Configuration Process.”](#)
- Configure preferred domain ID, FC ID persistence, and principle switch priority. See [Chapter 31, “Configuring Domain Parameters.”](#)
- Sophisticated SPAN diagnostics—The Cisco MDS 9000 Family provides industry-first intelligent diagnostics, protocol, decoding, and network analysis tools as well as integrated call-home capability for added reliability, faster problem resolution, and reduced service costs. See [Chapter 38, “Monitoring Network Traffic Using SPAN.”](#)
- Configure R_A_TOV, E_D_TOV. See the [“Fibre Channel Time Out Values”](#) section on page 39-2.
- Perform maintenance tasks for the director including maintaining firmware levels, accessing the director logs, and collecting data to support failure analysis. See [Chapter 41, “Monitoring System Processes and Logs.”](#)
- Display and clear port-level incident alerts. [“Clearing RLIR Information”](#) section on page 27-48.

Send documentation comments to mdsfeedback-doc@cisco.com.

FICON Port Numbering

With reference to the FICON feature, ports in Cisco MDS switches are identified by a statically defined 8-bit value known as the *port number*. Port numbers are assigned based on the module and the slot in the chassis. Port numbers cannot be changed and the first port in a switch always starts with a 0 (see [Figure 27-3](#)).

Figure 27-3 Port Number in the Cisco MDS 9000 Family



The FICON port number is assigned based on the front panel location of the port and is specific to the slot in which the module resides. Even if the module is a 16-port module, 32 port numbers are assigned to that module—regardless of the the module's physical presence in the chassis or the port status (up or down).



Note

Only Fibre Channel, PortChannel, and FCIP ports are mapped to FICON port numbers. Other types of interfaces do not have a corresponding port number.

[Table 27-1](#) lists the port number assignment for the Cisco MDS 9000 Family of switches and directors.

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 27-1 *FICON Port Numbering in the Cisco MDS 9000 Family*

| Product | Slot Number | Implemented Port Allocation | | Unimplemented Ports | Notes |
|-------------------------|-------------|-----------------------------|---------------------|------------------------------|--|
| | | To Ports | To PortChannel/FCIP | | |
| Cisco MDS 9200 Series | Slot 1 | 0 through 31 | 64 through 89 | 90 through 253 and port 255 | Similar to a switching module. |
| | Slot 2 | 32 through 63 | | | The first 16 port numbers in a 16-port module are used and the rest remain unused. |
| Cisco MDS 9506 Director | Slot 1 | 0 through 31 | 128 through 153 | 154 through 253 and port 255 | |
| | Slot 2 | 32 through 63 | | | |
| | Slot 3 | 64 through 95 | | | |
| | Slot 4 | 96 through 127 | | | |
| | Slot 5 | None | | | Supervisor modules are not allocated port numbers. |
| | Slot 6 | None | | | |
| Cisco MDS 9509 Director | Slot 1 | 0 through 31 | 224 through 249 | 250 through 253 and port 255 | The first 16 port numbers in a 16-port module are used and the rest remain unused. |
| | Slot 2 | 32 through 63 | | | |
| | Slot 3 | 64 through 95 | | | |
| | Slot 4 | 96 through 127 | | | |
| | Slot 5 | None | | | Supervisor modules are not allocated port numbers. |
| | Slot 6 | None | | | |
| | Slot 7 | 128 through 159 | | | The first 16 port numbers in a 16-port module are used and the rest remain unused. |
| | Slot 8 | 160 through 191 | | | |
| | Slot 9 | 192 through 223 | | | |

Port Addresses

By default, port numbers are the same as port addresses (see the “[Port Swapping](#)” section on [page 27-25](#)).

You can swap the port addresses by issuing the **ficon swap portnumber** command.

Implemented and Unimplemented Port Addresses

An implemented port refers to any port address that is available in the chassis (see [Table 27-1](#)).

An unimplemented port refers to any port address that is not available in the chassis (see [Table 27-1](#)).



Tip

An unimplemented port is prohibited from communicating with an implemented port in a FICON setup and cannot be configured.

Send documentation comments to mdsfeedback-doc@cisco.com.

Installed and Uninstalled Ports

An installed port refers to a port for which all required hardware is present. A specified port number in a VSAN can be implemented, and yet not installed, if any of the following conditions apply:

- The module is not present—for example, if module 1 is not physically present in slot 1 in a Cisco MDS 9509 Director, ports 0 to 31 are considered uninstalled.
- The small form-factor pluggable (SFP) port is not present—for example, if a 16-port module is inserted in slot 2 in a Cisco MDS 9509 Director, ports 48 to 63 are considered uninstalled.
- The port is not in a FICON-enabled VSAN—for example, if port 4 (of a 16-port module in slot 1) is configured in FICON-enabled VSAN 2, then only port 4 is installed and ports 0 to 3 and 5 to 15 are uninstalled—even if they are implemented in VSAN 2.

Another scenario is if VSANs 1 through 5 are FICON-enabled, and trunking-enabled interface fc1/1 has VSANs 3 through 10, then port address 0 is uninstalled in VSAN 1 and 2.

- The port is part of a PortChannel—for example, if interface fc 1/1 is part of PortChannel 5, port address 0 is uninstalled in all FICON VSANs. See [Table 27-1](#).

FICON Port Numbering Guidelines

The following guidelines apply to FICON port numbers:

- Supervisor modules do not have port number assignments.
- Port numbers are VSAN independent and do not change based on VSANs or TE ports.
- Each PortChannel must be explicitly associated with a FICON port number.
- When the port number for a physical PortChannel becomes uninstalled, the relevant PortChannel configuration is applied to the physical port.
- Each FCIP tunnel must be explicitly associated with a FICON port number. If the port numbers are not assigned for PortChannels or for FCIP tunnels, the associated ports will not come up.

See the [“FCIP and PortChannel Port Numbers”](#) section on page 27-9.

FCIP and PortChannel Port Numbers

FCIP and PortChannels cannot be used in a FICON-enabled VSAN unless they are explicitly bound to a port number.

See the [“Binding Port Numbers to PortChannels”](#) section on page 27-20 and the [“Binding Port Numbers to FCIP Interfaces”](#) section on page 27-20.

To find the first available port number to bind a FCIP or PortChannel interface use the **show ficon first-available port-number** command (see [Example 27-3](#)).



Tip

The **show ficon vsan portaddress brief** command displays the port number to interface mapping. You can assign port numbers in the PortChannel/FCIP range which are not already assigned to a PortChannel or FCIP interface (see [Example 27-4](#)).

Send documentation comments to mdsfeedback-doc@cisco.com.

FC ID Allocation

FICON requires a predictable and static FC ID allocation scheme. When FICON is enabled, the FC ID allocated to a device is based on the port address of the port to which it is attached. The port address forms the middle byte of the fabric address. Additionally, the last byte of the fabric address should be the same for all devices in the fabric. By default, the last byte value is 0 and can be configured (see the “FC ID Last Byte” section on page 27-16).




Note

You cannot configure persistent FC IDs in FICON-enabled VSANs.

Cisco MDS switches have a dynamic FC ID allocation scheme. When FICON is enabled or disabled on a VSAN, all the ports are flapped to switch from the dynamic to static FC IDs and vice versa (see Figure 27-4).

Figure 27-4 Static FC ID Allocation for FICON



113134

FICON Cascading

The Cisco MDS SAN-OS software allows multiple switches in a FICON network. To configure multiple switches, you must enable and configure fabric binding in that switch (see the “Fabric Binding Configuration” section on page 27-37).

Send documentation comments to mdsfeedback-doc@cisco.com.

FICON VSAN Prerequisites

To ensure that a FICON VSAN is operationally up, be sure to verify the following requirements:

- Set the default zone to permit, if you are not using the zoning feature. See the [“The Default Zone” section on page 15-11](#).
- Enable in-order delivery on the VSAN. See the [“In-Order Delivery” section on page 24-11](#).
- Enable (and if required, configure) fabric binding on the VSAN. See the [“Fabric Binding Configuration” section on page 27-37](#).
- Verify that conflicting persistent FC IDs do not exist in the switch. See [Chapter 31, “Configuring Domain Parameters.”](#)
- Verify that the configured domain ID and requested domain ID match. See [Chapter 31, “Configuring Domain Parameters.”](#)
- Add the CUP (area FE) to the zone, if you are using zoning. See the [“CUP In-Band Management” section on page 27-27](#).

If any of these requirements are not met, the FICON feature cannot be enabled.

Enabling FICON

By default FICON is disabled in all switches in the Cisco MDS 9000 Family.

You can enable FICON on a per VSAN basis in one of three ways:

- By using the automated **setup ficon** command.
See the [“Setting Up a Basic FICON Configuration” section on page 27-12](#) section in this chapter.
- Manually addressing each prerequisite.
See the [“Manually Enabling FICON” section on page 27-15](#).
- By using the Device Manager (refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*).

Effects of Enabling FICON

When you enable the FICON feature in Cisco MDS switches, the following apply:

- You cannot disable in-order delivery for the FICON-enabled VSAN.
- You cannot disable fabric binding or static domain ID configurations for the FICON-enabled VSAN.
- The load balancing scheme is changed to Source ID (SID)—Destination ID (DID). You cannot change it back to SID—DID—OXID.
- The IPL configuration file is automatically created.
See the [“FICON Configuration Files” section on page 27-23](#).

Send documentation comments to mdsfeedback-doc@cisco.com.

Setting Up a Basic FICON Configuration

This section steps you through the procedure to set up FICON on a specified VSAN in a Cisco MDS 9000 Family switch.



Note

Press **Ctrl-C** at any prompt to skip the remaining configuration options and proceed with what is configured until that point.



Tip

If you do not wish to answer a previously configured question, or if you wish to skip answers to any questions, press **Enter**. If a default answer is not available (for example, switch name), the switch uses what was previously configured and skips to the next question.

To enable and set up FICON, follow these steps.

- Step 1** Issue the **setup ficon** command at the EXEC command mode.

```
switch# setup ficon
      --- Ficon Configuration Dialog ---
```

This setup utility will guide you through basic Ficon Configuration on the system.

Press Enter if you want to skip any dialog. Use ctrl-c at anytime to skip all remaining dialogs.

- Step 2** Enter **yes** (the default is **yes**) to enter the basic FICON configuration setup.

```
Would you like to enter the basic configuration dialog (yes/no) [yes]: yes
```

The FICON setup utility guides you through the basic configuration process. Press **Ctrl-C** at any prompt to end the configuration process.

- Step 3** Enter the VSAN number for which FICON should be enabled.

```
Enter vsan [1-4093]: 2
```

- Step 4** Enter **yes** (the default is **yes**) to create a new VSAN.

```
vsan 2 does not exist, create it? (yes/no) [yes]: yes
```

- Step 5** Enter **yes** (the default is **yes**) to confirm your VSAN choice:

```
Enable ficon on this vsan? (yes/no) [yes]: yes
```



Note

At this point, the software creates the VSAN if it does not already exist.

- Step 6** Enter the domain ID number for the specified FICON VSAN.

```
Configure domain-id for this ficon vsan (1-239): 2
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Step 7 Enter **yes** (the default is **no**) to set up FICON in cascaded mode. If you enter **no**, skip to [Step 8](#) (see the “[CUP In-Band Management](#)” section in this chapter).

Would you like to configure ficon in cascaded mode: (yes/no) [no]: **yes**

a. Assign the peer WWN for the FICON: CUP.

Configure peer wwn (hh:hh:hh:hh:hh:hh:hh:hh): **11:00:02:01:aa:bb:cc:00**

b. Assign the peer domain ID for the FICON: CUP

Configure peer domain (1-239) :**4**

c. Enter **yes** if you wish to configure additional peers (and repeat Steps [7a](#) and [7b](#)). Enter **no**, if you do wish to configure additional peers.

Would you like to configure additional peers: (yes/no) [no]: **no**

Step 8 Enter **yes** (the default is **yes**) to deny SNMP permission to modify existing port connectivity parameters (see the “[FICON SNMP Control](#)” section in this chapter).

Enable SNMP to modify port connectivity parameters? (yes/no) [yes]: **yes**

Step 9 Enter **no** (the default is **no**) to disable the host (mainframe) to modify the port connectivity parameters, if required (see the “[FICON Host Control](#)” section in this chapter.).

Disable Host from modifying port connectivity parameters? (yes/no) [no]: **no**

Step 10 Enter **yes** (the default is **yes**) to enable the **active equals saved** feature (see the “[Running Configuration Automatic Save](#)” section in this chapter.).

Enable active=saved? (yes/no) [yes]: **yes**

Step 11 Enter **yes** (the default is **yes**) if you wish to configure additional FICON VSANs.

Would you like to configure additional ficon vsans (yes/no) [yes]: **yes**

Step 12 Review and edit the configuration that you have just entered.

Step 13 Enter **no** (the default is **no**) if you are satisfied with the configuration.



Note

For documentation purposes, the following configuration displays three VSANs with different FICON settings. These settings provide a sample output for different FICON scenarios.

The following configuration will be applied:

```
fcdomain domain 2 static vsan 1
fcdomain restart disruptive vsan 1
fabric-binding database vsan 1
swwn 11:00:02:01:aa:bb:cc:00 domain 4
fabric-binding activate vsan 1
zone default-zone permit vsan 1
ficon vsan 1
no host port control

fcdomain domain 3 static vsan 2
fcdomain restart disruptive vsan 2
fabric-binding activate vsan 2 force
zone default-zone permit vsan 2
ficon vsan 2
no host port control
no active equals saved
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```

vsan database
vsan 3
fcdomain domain 5 static vsan 3
fcdomain restart disruptive vsan 3
fabric-binding activate vsan 3 force
zone default-zone permit vsan 3
ficon vsan 3
no snmp port control
no active equals saved

```

Would you like to edit the configuration? (yes/no) [no]: **no**

Step 14 Enter **yes** (the default is **yes**) to use and save this configuration. The implemented commands are displayed. After FICON is enabled for the specified VSAN, you are returned to the EXEC mode switch prompt.

Use this configuration and apply it? (yes/no) [yes]: **yes**

```

`fcdomain domain 2 static vsan 1`
`fcdomain restart disruptive vsan 1`
`fabric-binding database vsan 1`
`swmn 11:00:02:01:aa:bb:cc:00 domain 4`
`fabric-binding activate vsan 1`
`zone default-zone permit vsan 1`
`ficon vsan 1`
`no host port control`

`fcdomain domain 3 static vsan 2`
`fcdomain restart disruptive vsan 2`
`fabric-binding activate vsan 2 force`
`zone default-zone permit vsan 2`
`ficon vsan 2`
`no host port control`
`no active equals saved`

```



Note If a new VSAN is created, two additional commands are displayed— **vsan database** and **vsan number**.

```

`vsan database`
`vsan 3`
`in-order-guarantee vsan 3`
`fcdomain domain 2 static vsan 3`
`fcdomain restart disruptive vsan 3`
`fabric-binding activate vsan 3 force`
`zone default-zone permit vsan 3`
`ficon vsan 3`
`no snmp port control`
Performing fast copy config...done.
switch#

```

Send documentation comments to mdsfeedback-doc@cisco.com.

Manually Enabling FICON



Tip

This section describes the procedure to manually enable FICON on a VSAN. If you have already enabled FICON on the required VSAN using the automated setup (recommended), skip to the [“Running Configuration Automatic Save”](#) section on page 27-19.

To manually enable FICON on a VSAN, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# vsan database switch(config-vsan-db)# vsan 5 switch(config-vsan-db)# do show vsan usage 4 vsan configured configured vsans:1-2,5,26 vsans available for configuration:3-4,6-25,27-4093 switch(config-vsan-db)# exit | Enables VSAN 5. |
| Step 3 | switch(config)# in-order-guarantee vsan 5 | Activates in-order delivery for VSAN 5. See Chapter 24, “Configuring Fibre Channel Routing Services and Protocols.” |
| Step 4 | switch(config)# fcdomain domain 2 static vsan 2 | Configures the domain ID for VSAN 2. See Chapter 31, “Configuring Domain Parameters.” |
| Step 5 | switch(config)# fabric-binding activate vsan 2 force | Activates fabric binding on VSAN 2. See the “Fabric Binding Configuration” section in this chapter. |
| Step 6 | switch(config)# zone default-zone permit vsan 2 | Sets the default zone to permit for VSAN 2. See the “CUP In-Band Management” section in this chapter. |
| Step 7 | switch(config)# ficon vsan 2 switch(config-ficon)# | Enables FICON on VSAN 2. |
| | switch(config)# no ficon vsan 6 | Disables the FICON feature on VSAN 6. |
| Step 8 | switch(config-ficon)# no host port control | Prohibits mainframe users from moving the switch to an offline state. See the “Host Moves the Switch Offline” section in this chapter. |

Send documentation comments to mdsfeedback-doc@cisco.com.

The code-page Option

FICON strings are coded in Extended Binary-Coded Decimal Interchange Code (EBCDIC) format. Refer to your mainframe documentation for details on the code page options.

Cisco MDS switches support **international-5**, **france**, **brazil**, **germany**, **italy**, **japan**, **spain-latinamerica**, **uk**, and **us-canada** (default) EBCDIC format options.



Tip

This is an optional configuration. If you are not sure of the EBCDIC format to be used, we recommend retaining the **us-canada** (default) option.

To configure the **code-page** option in a VSAN, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ficon vsan 2 switch(config-ficon)# | Enables FICON on VSAN 2. |
| Step 3 | switch(config-ficon)# code-page italy switch(config-ficon)# no code-page | Configures the italy EBCDIC format. Reverts to the factory default of using the us-canada EBCDIC format. |

FC ID Last Byte



Caution

If the FICON feature is configured in cascaded mode, the Cisco MDS Switches use ISLs to connect to other switches.

FICON requires the last byte of the fabric address to be the same for all allocated FC IDs. By default, this value is set to 0. You can only change the FC ID last byte when the FICON switch is in the offline state.

See the [“Moving a FICON VSAN to an Offline State” section on page 27-27](#).

To assign the last byte for the FC ID, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ficon vsan 2 switch(config-ficon)# | Enables FICON on VSAN 2. |
| Step 3 | switch(config-ficon)# fcid-last-byte 12 switch(config-ficon)# no fcid-last-byte 3 | Assigns the last byte FC ID for the fabric address. Removes the configured last byte FC ID for the fabric address and reverts to the factory default of 0. |

Send documentation comments to mdsfeedback-doc@cisco.com.

FICON Host Control

The commands included in this section allow the host (mainframe) to control the Cisco MDS switch.

Host Moves the Switch Offline

By default, hosts are allowed to move the switch to an offline state.

Use the **host control switch offline** command to allow the host to move the switch to an offline state.

To allow the host to move the switch to an offline state, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ficon vsan 2 switch(config-ficon)# | Enables FICON on VSAN 2. |
| Step 3 | switch(config-ficon)# no host control switch offline | Prohibits mainframe users from moving the switch to an offline state. |
| | switch(config-ficon)# host control switch offline | Allows the host to move the switch to an offline state (default) and shuts down the ports. |

Host Changes FICON Port Parameters

By default, mainframe users are not allowed to configure FICON parameters on Cisco MDS switches—they can only query the switch.

Use the **host port control** command to permit mainframe users to configure FICON parameters.

To configure mainframe access, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ficon vsan 2 switch(config-ficon)# | Enables FICON on VSAN 2. |
| Step 3 | switch(config-ficon)# no host port control | Prohibits mainframe users from configuring FICON parameters on the Cisco MDS switch. |
| | switch(config-ficon)# host port control | Allows mainframe users to configure FICON parameters on the Cisco MDS switch (default). |

Host Controls the Time Stamp

By default, the clock in each VSAN is the same as the switch hardware clock. Each VSAN in a Cisco MDS 9000 Family switch represents a virtual director. The clock and time present in each virtual director can be different. To maintain separate clocks for each VSAN, the Cisco SAN-OS software maintains the difference of the VSAN-specific clock and the hardware-based director clock. When a host (mainframe) sets the time, the Cisco SAN-OS software updates this difference between the clocks. When a host reads the clock, it computes the difference between the VSAN-clock and the current director hardware clock and presents a value to the mainframe.

The VSAN-clock's current time is reported in the output of **show ficon vsan vsan-id**, **show ficon**, and **show accounting log** commands.

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure host control, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ficon vsan 2 switch(config-ficon)# | Enables FICON on VSAN 2. |
| Step 3 | switch(config-ficon)# no host set-timestamp | Prohibits mainframe users from changing the VSAN-specific clock. |
| | switch(config-ficon)# host set-timestamp | Allows the host to set the clock on this switch (default). |

Time Stamp Cleanup



Note

You can clear time stamps only from the Cisco MDS switch—not the mainframe.

Use the **clear ficon vsan vsan-id timestamp** command in EXEC mode to clear the VSAN clock.

```
switch# clear ficon vsan 20 timestamp
```

FICON SNMP Control

By default, SNMP users can configure FICON parameters through the Cisco MDS 9000 Family Fabric Manager.



Note

If you disable SNMP use in the Cisco MDS switch, you cannot configure FICON parameters using the Fabric Manager.

You can prohibit this access, if required, by issuing the **no snmp port control** command.

To configure SNMP control, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ficon vsan 2 switch(config-ficon)# | Enables FICON on VSAN 2. |
| Step 3 | switch(config-ficon)# no snmp port control | Prohibits SNMP users from configuring FICON parameters. |
| | switch(config-ficon)# snmp port control | Allows SNMP users to configure FICON parameters (default). |

Send documentation comments to mdsfeedback-doc@cisco.com.

Running Configuration Automatic Save

Table 27-2 displays the results of **active equals saved** command and the implicit **copy running start** command in various scenarios.

If **active equals saved** is enabled in any FICON-enabled VSAN in the fabric, then the following apply (see Number 1 and 2 in Table 27-2):

- All configuration changes (FICON-specific or not) are automatically saved to persistent storage (implicit **copy running start**) and stored in the startup configuration.
- FICON-specific configuration changes are immediately saved to the IPL file (see the “[FICON Configuration Files](#)” section on page 27-23).

If **active equals saved** is not enabled in any FICON-enabled VSAN in the fabric, then FICON-specific configuration changes are not saved in the IPL file and an implicit **copy running start** is not issued—you must issue the **copy running start** command explicitly (see Number 3 in Table 27-2):

Table 27-2 Saving the Active FICON and Switch Configuration

| Number | FICON-enabled VSAN? | active equals saved Enabled? | Implicit ¹ copy running start Issued? | Notes |
|--------|---------------------|------------------------------|--|---|
| 1 | Yes | Yes (in all FICON VSANs) | Implicit | FICON changes written to the IPL file. Non-FICON changes saved to startup configuration and persistent storage. |
| 2 | | Yes (even in one FICON VSAN) | Implicit | FICON changes written to IPL file for only the VSAN which has active equals saved enabled. Non-FICON changes saved to startup configuration and persistent storage. |
| 3 | | Not in any FICON VSAN | Not implicit | FICON changes are not written to the IPL file. Non-FICON changes are saved in persistent storage—only if you explicitly issue the copy running start command. |
| 4 | No | Not applicable | | |

1. When the Cisco SAN-OS software implicitly issues a **copy running start** command in the Cisco MDS switch, only a binary configuration is generated—an ASCII configuration is not generated (see [Example 27-16](#)). If you wish to generate an additional ASCII configuration at this stage, you must explicitly issue the **copy running start** command again.



Note

If **active equals saved** is enabled, the Cisco SAN-OS software ensures that you do not have to perform the **copy running startup** command for the FICON configuration as well. If your switch or fabric consists of multiple FICON-enabled VSANs, and one of these VSANs have **active equals saved** enabled, changes made to the non-FICON configuration results in all configurations being saved to the startup configuration.

Send documentation comments to mdsfeedback-doc@cisco.com.

To automatically save the running configuration, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ficon vsan 2 switch(config-ficon)# | Enables FICON on VSAN 2. |
| Step 3 | switch(config-ficon)# active equals saved | Enables the automatic save feature for all VSANs in the switch or fabric. |
| | switch(config-ficon)# no active equals saved | Disables automatic save for this VSAN. |

Binding Port Numbers to PortChannels



Caution

All port number assignments to PortChannels/FCIP interfaces are lost (cannot be retrieved) when FICON is disabled on all VSANs.

You can bind (or associate) a PortChannel with a FICON port number to bring up that interface.

To bind a PortChannel with a FICON port number, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# interface Port-channel 1 switch(config-if)# | Enters the PortChannel interface configuration mode. |
| Step 3 | switch(config-if)# ficon portnumber 234 | Assigns the FICON port number to the selected PortChannel port. |

Binding Port Numbers to FCIP Interfaces

You can bind (or associate) a FCIP interface with a FICON port number to bring up that interface.

To bind a FCIP interface with a FICON port number, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch1(config)# interface fcip 51 switch1(config-if)# | Creates a FCIP interface (51). |
| Step 3 | switch(config-if)# ficon portnumber 208 | Assigns the FICON port number to the selected FCIP interface. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Configuring FICON Ports

You can perform FICON configurations on a per-port address basis in the Cisco MDS 9000 Family of switches.

Even if a port is uninstalled, the port address-based configuration is accepted by the Cisco MDS switch. This configuration is applied to the port when the port becomes installed.

Port Blocking

If you block a port, the port is retained in the operationally down state. If you unblock a port, a port initialization is attempted. When a port is blocked, data and control traffic are not allowed on that port.

Physical Fibre Channel port blocks will continue to transmit an Off-Line State (OLS) primitive sequence on a blocked port.



Caution

You cannot block or prohibit the CUP port (0XFE).

If a port is shut down, unblocking that port does not initialize the port.



Note

The **shutdown/no shutdown** port state is independent of the **block/no block** port state.

To block or unblock port addresses in a VSAN, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ficon vsan 2 switch(config-ficon)# | Enables FICON on VSAN 2. |
| Step 3 | switch(config-ficon)# portaddress 1 - 5 switch(config-ficon-portaddr)# | Selects port address 1 to 5 for further configuration. |
| Step 4 | switch(config-ficon-portaddr)# block | Disables a range of port addresses and retains it in the operationally down state. |
| | switch(config-ficon-portaddr)# no block | Enables the selected port address and reverts to the factory default of the port address not being blocked. |

Port Prohibiting

To prevent implemented ports from talking to each other, you can configure prohibits between two or more ports. If you prohibit ports, the specified ports are prevented from communicating with each other.



Note

Unimplemented ports are always prohibited.



Tip

You cannot prohibit a PortChannel or FCIP interface.

Send documentation comments to mdsfeedback-doc@cisco.com.

Prohibit configurations are always symmetrically applied—if you prohibit Port 0 from talking to port 15, port 15 is automatically prohibited from talking to port 0.



Note

If an interface is already configured in E or TE mode and you try to prohibit that port, your prohibit configuration is rejected. Similarly, if a port is not up and you prohibit that port, the port is not allowed to come up in E mode nor in TE mode.

To prohibit port addresses in a VSAN, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ficon vsan 2 switch(config-ficon)# | Enables FICON on VSAN 2. |
| Step 3 | switch(config-ficon)# portaddress 7 switch(config-ficon-portaddr)# | Selects port address 7 for further configuration. |
| Step 4 | switch(config-ficon-portaddr)# prohibit portaddress 3-5 | Prohibits port address 7 in VSAN 2 from talking to ports 3, 4, and 5. |
| | switch(config-ficon-portaddr)# no prohibit portaddress 5 | Removes port address 5 from a previously prohibited state. |

Port Address Name Assignment

To assign a port address name, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ficon vsan 2 switch(config-ficon)# | Enables FICON on VSAN 2. |
| Step 3 | switch(config-ficon)# portaddress 7 switch(config-ficon-portaddr)# | Selects port address 7 for further configuration. |
| Step 4 | switch(config-ficon-portaddr)# name SampleName | Assigns a name to the port address. |
| | switch(config-ficon-portaddr)# no name SampleName | Deletes a previously configured port address name. |

Note The port address name is restricted to 24 alphanumeric characters.

Send documentation comments to mdsfeedback-doc@cisco.com.

FICON Configuration Files

You can save up to 16 FICON configuration files on each FICON-enabled VSAN (in persistent storage). The file format is proprietary to IBM. These files can be read and written by IBM hosts using the in-band CUP protocol. Additionally, you can use the Cisco MDS CLI or Fabric Manager applications to operate these FICON configuration files.

**Note**

Multiple FICON configuration files with the same name can exist in the same switch, provided they reside in different VSANs. For example, you can create a configuration file named XYZ in both VSAN 1 and VSAN 3.

When you enable the FICON feature in a VSAN, the switches always use the startup FICON configuration file, called IPL. This file is created with a default configuration as soon as FICON is enabled in a VSAN.

**Caution**

When FICON is disabled on a VSAN, all the FICON configuration files are irretrievably lost.

FICON configuration files contain the following configuration for each implemented port address:

- Block
- Prohibit mask
- Port address name

**Note**

Normal configuration files used by Cisco MDS switches include FICON-enabled attributes for a VSAN, port number mapping for PortChannels and FCIP interfaces, port number to port address mapping, port and trunk allowed VSAN configuration for ports, in-order guarantee, configuring static domain ID, and fabric binding configuration.

See the [“Working with Configuration Files” section on page 4-24](#) for details on the normal configuration files used by Cisco MDS switches.

Accessing FICON Configuration Files

Only one user can access the configuration file at any given time:

- If this file is being accessed by user 1, user 2 cannot access this file.
- If user 2 does attempt to access this file, an error is issued to user 2.
- If user 1 is inactive for more than 15 seconds, the file is automatically closed and available for use by any other permitted user.

FICON configuration files can be accessed by any host, SNMP, or CLI user who is permitted to access the switch. The locking mechanism in the Cisco SAN-OS software restricts access to one user at a time per file. This lock applies to newly created files and previously saved files. Before accessing any file, you must lock the file and obtain the file key. A new file key is used by the locking mechanism for each lock request. The key is discarded when the lock timeout of 15 seconds expires. The lock timeout value cannot be changed.

Send documentation comments to mdsfeedback-doc@cisco.com.

Applying the FICON Configuration Files

The configuration from the saved files can be applied to the running configuration by using the **ficon vsan number apply file filename** command. For example:

```
switch# ficon vsan 2 apply file SampleFile
```

Editing FICON Configuration Files

The configuration file submode allows you to create and edit FICON configuration files. If a specified file does not exist, it is created. Up to 16 files can be saved. Each file name is restricted to eight alphanumeric characters.

To edit the contents of a specified FICON configuration file, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ficon vsan 2 switch(config-ficon)# | Enables FICON on VSAN 2. |
| Step 3 | switch(config-ficon)# file IplFile1 switch(config-ficon-file)# | Accesses the FICON configuration file called IplFile1 for VSAN 2. If this file does not exist, it is created. Note All FICON file names are restricted to eight alphanumeric characters. |
| | switch(config-ficon)# no file IplFileA | Deletes a previously created FICON configuration file. |
| Step 4 | switch(config-ficon-file)# portaddress 3 switch(config-ficon-file-portaddr)# | Enters the submode for port address 3 to edit the contents of the configuration file named IplFile1. Note The running configuration is not applied to the current configuration. The configuration is only applied when the ficon vsan number apply file filename command is issued. |
| Step 5 | switch(config-ficon-file-portaddr)# prohibit portaddress 5 | Edits the content of the configuration file named IplFile1 by prohibiting port address 5 from accessing port address 3. |
| Step 6 | switch(config-ficon-file-portaddr)# block | Edits the content of the configuration file named IplFile1 by blocking a range of port addresses and retaining them in the operationally down state. |
| Step 7 | switch(config-ficon-file-portaddr)# name P3 | Edits the content of the configuration file named IplFile1 by assigning the name P3 to port address 3. If the name did not exist, it is created. If it existed, it is overwritten. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Copying FICON Configuration Files

Use the **ficon vsan vsan-id copy file exiting-file-name save-as-file-name** command in EXEC mode to copy an existing FICON configuration file.

```
switch# ficon vsan 20 copy file IPL IPL3
```

You can see the list of existing configuration files by issuing the **show ficon vsan vsan-id** command.

```
switch# show ficon vsan 20
Ficon information for VSAN 20
  Ficon is online
  VSAN is active
  Host port control is Enabled
  Host offline control is Enabled
  User alert mode is Disabled
  SNMP port control is Enabled
  Host set director timestamp is Enabled
  Active=Saved is Enabled
  Number of implemented ports are 250
  Key Counter is 5
  FCID last byte is 0
  Date/Time is same as system time (Wed Dec 3 20:10:45.924591 2003)
  Device Allegiance not locked
  Codepage is us-canada
Saved configuration files
  IPL
  IPL3
```

Port Swapping

The FICON port swap feature is only provided for maintenance purposes.

The FICON port swapping feature causes all configuration associated with *old-port-number* and *new-port-number* to be swapped, including VSAN configurations.

Cisco MDS switches allow port swapping for non-existent ports as follows:

- Only FICON-specific configurations (prohibit, block, and port address mapping) are swapped.
- No other system configuration is swapped.
- All other system configurations are only maintained for existing ports.



Tip

If **active equals saved** is enabled on any FICON VSAN, then the swapped configuration is automatically saved to startup. Otherwise, you must explicitly save the running configuration immediately after swapping the ports.

Once you swap ports, the switch automatically performs the following actions:

- Shuts down both the old and new ports.
- Swaps the port configuration.
- If you attempt to bring the port up, you must explicitly shut down the port to resume traffic.

The **ficon swap portnumber** command is only associated with the two ports concerned. You must issue this VSAN-independent command from EXEC mode.

Send documentation comments to mdsfeedback-doc@cisco.com.

If you attempt to bring the port up by specifying the **ficon swap portnumber** *old-port-number new-port-number* **after swap noshut** command, you must explicitly issue the **no shutdown** command to resume traffic.

To swap physical Fibre Channel ports, follow these steps:

-
- Step 1** Issue the **ficon swap portnumber** *old-port-number new-port-number* command in EXEC mode. The specified ports are operationally shut down.
- Step 2** Physically swap the front panel port cables between the two ports.
- Step 3** Issue the **no shutdown** command on each port to enable traffic flow.



Note If you specify the **ficon swap portnumber** *old-port-number new-port-number* **after swap noshut** command, the ports are automatically initialized.

Port Swapping Guidelines

Be sure to follow these guidelines when using the FICON port swap feature:

- Port swapping is not supported for logical ports (PortChannels, FCIP links). Neither the *old-port-number* nor the *new-port-number* can be a logical port.
- Port swapping is not supported between physical ports that are part of a PortChannel. Neither the *old-port-number* nor the *new-port-number* can be a physical port that is part of a PortChannel.
- Before performing a port swap, the Cisco SAN-OS software performs a compatibility check. If the two ports have incompatible configurations, the port swap is rejected with an appropriate reason code. For example, if a port with BB_credits as 25 is being swapped with an OSM port for which a maximum of 12 BB_credits is allowed (not a configurable parameter), the port swapping operation is rejected.
- Before performing a port swap, the Cisco SAN-OS software performs a compatibility check to verify the extended BB_credits configuration.
- If ports have default values (for some incompatible parameters), then a port swap operation is allowed and the ports retain their default values.
- Port tracking information is not included in port swapping. This information must be configured separately (see [Chapter 33, “Tracking and Redirecting Traffic”](#)).



Note

The 32-port module guidelines also apply for port swapping configurations (see the [“32-Port Configuration Guidelines”](#) section on page 12-8).

Send documentation comments to mdsfeedback-doc@cisco.com.

Moving a FICON VSAN to an Offline State

Use the EXEC-level **ficon vsan vsan-id offline** command to log out all ports in the VSAN that needs to be suspended.

Use the EXEC-level **ficon vsan vsan-id online** command to remove the offline condition and to allow ports to log on again.



Note

This command can be issued by the host if the host is allowed to do so (see the “[Host Moves the Switch Offline](#)” section in this chapter).

Clearing FICON Device Allegiance

FICON requires serialization of access among multiple mainframes, CLI, and SNMP sessions be maintained on Cisco MDS 9000 Family switches by controlling device allegiance for the currently executing session. Any other session is denied permission to perform configuration changes unless the required allegiance is available.



Caution

This task aborts the currently executing session.

You can clear the current device allegiance by issuing the **clear ficon vsan vsan-id allegiance** command in EXEC mode.

```
switch# clear ficon vsan 1 allegiance
```

CUP In-Band Management

The Control Unit Port (CUP) protocol configures access control and provides unified storage management capabilities from a mainframe computer. Cisco MDS 9000 FICON-enabled switches are fully IBM CUP standard compliant for in-band management using the IBM S/A OS/390 I/O operations console.



Note

The CUP specification is proprietary to IBM.

CUP is supported by switches and directors in the Cisco MDS 9000 Family. The CUP function allows the mainframe to manage the Cisco MDS switches.

Host communication includes control functions such as blocking and unblocking ports, as well as monitoring and error reporting functions.

Send documentation comments to mdsfeedback-doc@cisco.com.

Placing CUPs in a Zone

To place the CUP in a zone, follow these steps.

- Step 1** Set the default zone to permit for the required VSAN.

```
switch# config t
switch(config)# zone default-zone permit vsan 20
```

- Step 2** Issue the **show fcns database** command for the required VSAN and obtain the required FICON CUP WWN.

```
switch# show fcns database vsan 20
```

VSAN 20:

| FCID | TYPE | PWWN | (VENDOR) | FC4-TYPE:FEATURE |
|----------|------|--------------------------------|----------|-----------------------|
| 0x0d0d00 | N | 50:06:04:88:00:1d:60:83 | (EMC) | FICON:CU |
| 0x0dfe00 | N | 25:00:00:0c:ce:5c:5e:c2 | (Cisco) | FICON:CUP |
| 0x200400 | N | 50:05:07:63:00:c2:82:d3 | (IBM) | scsi-fcp FICON:CU f.. |
| 0x200800 | N | 50:05:07:64:01:40:15:0f | (IBM) | FICON:CH |
| 0x20fe00 | N | 20:00:00:0c:30:ac:9e:82 | (Cisco) | FICON:CUP |

Total number of entries = 5



Note If more than one FICON:CUP WWN exists in this fabric, be sure to add all the FICON:CUP WWN PWWNs to the required zone. The previous example displays multiple FICON:CUP occurrences to indicate a cascade configuration.

- Step 3** Add the identified FICON:CUP WWN to the zone database.

```
switch(config)# zone name Zone1 vsan 20
switch(config-zone)# member pwwn 25:00:00:0c:ce:5c:5e:c2
```

Displaying FICON Information

Use the **show** commands to display all FICON information configured on this switch (see Examples 27-1 to 27-15).

Receiving FICON Alerts

In Example 27-1 the **user alert mode is enabled** output confirms that you will receive an alert to indicate any changes in the FICON configuration.

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 27-1 Displays Configured FICON Information

```
switch# show ficon
Ficon information for VSAN 20
  Ficon is online
  VSAN is active
  Host port control is Enabled
  Host offline control is Enabled
  User alert mode is Enabled
  SNMP port control is Enabled
  Host set director timestamp is Enabled
  Active=Saved is Disabled
  Number of implemented ports are 250
  Key Counter is 73723
  FCID last byte is 0
  Date/Time is set by host to Sun Jun 26 00:04:06.991999 1904
  Device allegiance is locked by Host
  Codepage is us-canada
  Saved configuration files
    IPL
    _TSIRN00
```

Displaying FICON Port Address Information

Examples 27-2 to 27-5 display FICON Port Address information.

Example 27-2 Displays Port Address Information

```
switch# show ficon vsan 2 portaddress
Port Address 1 is not installed in vsan 2
  Port number is 1, Interface is fc1/1
  Port name is
  Port is not admin blocked
  Prohibited port addresses are 0,241-253,255

Port Address 2 is not installed in vsan 2
  Port number is 2, Interface is fc1/2
  Port name is
  Port is not admin blocked
  Prohibited port addresses are 0,241-253,255
...
Port Address 249 is not installed in vsan 2
  Port name is
  Port is not admin blocked
  Prohibited port addresses are 0,241-253,255

Port Address 250 is not installed in vsan 2
  Port name is
  Port is not admin blocked
  Prohibited port addresses are 0,241-253,255
```

Example 27-3 Displays the Available Port Numbers

```
switch# show ficon first-available port-number
Port number 129(0x81) is available
```

In Example 27-4, the interface column is populated with the corresponding interface if the port number is installed. If the port number is uninstalled, this space remains blank and indicates an unbound port number. For example, 56 is an unbound port number in Example 27-4.

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 27-4 Displays Port Address Information in a Brief Format

```
switch# show ficon vsan 2 portaddress 50-55 brief
```

| Port Address | Port Number | Interface | Admin Blocked | Status | Oper Mode | FCID |
|--------------|-------------|-----------|---------------|--------------|-----------|----------|
| 50 | 50 | fc2/18 | on | fcotAbsent | -- | -- |
| 51 | 51 | fc2/19 | off | fcotAbsent | -- | -- |
| 52 | 52 | fc2/20 | off | fcotAbsent | -- | -- |
| 53 | 53 | fc2/21 | off | fcotAbsent | -- | -- |
| 54 | 54 | fc2/22 | off | notConnected | -- | -- |
| 55 | 55 | fc2/23 | off | up | FL | 0xea0000 |
| 56 | 56 | | off | up | FL | 0xea0000 |

Example 27-5 displays the counters in FICON version format 1 (32-bit format)

Example 27-5 Displays Port Address Counter Information

```
switch# show ficon vsan 20 portaddress 8 counters
Port Address 8(0x8) is up in vsan 20
  Port number is 8(0x8), Interface is fc1/8
  Version presented 1, Counter size 32b
  242811 frames input, 9912794 words
    484 class-2 frames, 242302 class-3 frames
    0 link control frames, 0 multicast frames
    0 disparity errors inside frames
    0 disparity errors outside frames
    0 frames too big, 0 frames too small
    0 crc errors, 0 eof errors
    0 invalid ordered sets
    0 frames discarded c3
    0 address id errors
  116620 frames output, 10609188 words
    0 frame pacing time
  0 link failures
  0 loss of sync
  0 loss of signal
  0 primitive seq prot errors
  0 invalid transmission words
  1 lrr input, 0 ols input, 5 ols output
  0 error summary
```

Displaying IPL File Information

Examples 27-6 to 27-5 display FICON Port Address information.

Example 27-6 Displays the Contents of the Specified FICON Configuration File

```
switch# show ficon vsan 3 file IPL
FICON configuration file IPL      in vsan 3
  Port address 1
    Port name is
    Port is not blocked
    Prohibited port addresses are 0,81-253,255

  Port address 2
    Port name is
    Port is not blocked
    Prohibited port addresses are 0,81-253,255
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
Port address 3
  Port name is
  Port is not blocked
  Prohibited port addresses are 0,81-253,255

Port address 4
  Port name is
  Port is not blocked
  Prohibited port addresses are 0,81-253,255

...
Port address 80
  Port name is
  Port is not blocked
  Prohibited port addresses are 0,81-253,255

Port address 254
  Port name is
  Port is not blocked
  Prohibited port addresses are 0,81-253,255
```

Example 27-7 Displays All FICON Configuration Files

```
switch# show ficon vsan 2
Ficon information for VSAN 2
  Ficon is enabled
  VSAN is active
  Host control is Enabled
  Host offline control is Enabled
  Clock alert mode is Disabled
  User alert mode is Disabled
  SNMP control is Disabled
  Active=Saved is Disabled
  Number of implemented ports are 250
  Key Counter is 9
  FCID last byte is 0
  Date/Time is same as system time(Sun Dec 14 01:26:30.273402 1980)
  Device Allegiance not locked
  Codepage is us-canada
Saved configuration files
  IPL
  IPLFILE1
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 27-8 *Displays the Specified Port Addresses for a FICON Configuration File*

```
switch# show ficon vsan 2 file iplfile1 portaddress 1-7
FICON configuration file IPLFILE1 in vsan 2
  Port address 1
    Port name is
    Port is not blocked
    Prohibited port addresses are 0,241-253,255

  Port address 2
    Port name is
    Port is not blocked
    Prohibited port addresses are 0,241-253,255

  Port address 3
    Port name is P3
    Port is not blocked
    Prohibited port addresses are 0,241-253,255
...
  Port address 7
    Port name is
    Port is not blocked
    Prohibited port addresses are 0,241-253,255
```

Displaying the Configured FICON State

If FICON is enabled on a VSAN, you can view the port address information for that VSAN (see [Example 27-9](#)).

Example 27-9 *Displays the Specified Port Address When FICON Is Enabled*

```
switch# show ficon vsan 2 portaddress 55
Port Address 55 is not installed in vsan 2
  Port number is 55, Interface is fc2/23
  Port name is
  Port is not admin blocked
  Prohibited port addresses are 0,241-253,255
  Admin port mode is FL
  Port mode is FL, FCID is 0xea0000
```


Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying a Ports Administrative State

Examples 27-10 to 27-11 display the administrative state of a FICON port. If the port is blocked, the **show ficon vsan number portaddress number** command displays the blocked state of the port. If a specific port is prohibited, this command also displays the specifically prohibited port (3) along with the ports that are prohibited by default (0, 241 to 253, and 255). If a name is assigned, that name is also displayed.

Example 27-10 Displays an Administratively Unblocked Port

```
switch# show ficon vsan 2 portaddress 2
Port Address 2(0x2) is not installed in vsan 2
  Port number is 2(0x2), Interface is fcl/2
  Port name is
Port is not admin blocked
  Prohibited port addresses are 0,241-253,255(0,0xf1-0xfd,0xff)
  Admin port mode is auto
  Peer is Unknown
```

Example 27-11 Displays an Administratively Blocked Port

```
switch# show ficon vsan 2 portaddress 1
Port Address 2(0x2) is not installed in vsan 2
  Port number is 2(0x2), Interface is fcl/2
  Port name is SampleName
Port is admin blocked
  Prohibited port addresses are 0,241-253,255(0,0xf1-0xfd,0xff)
  Admin port mode is auto
  Peer is Unknown
```

Displaying Control Unit Information

Example 27-12 displays configured control device information.

Example 27-12 Displays Control Unit Information

```
switch# show ficon control-device sb3
Control Unit Image:0x80b9c2c
VSAN:20 CU:0x20fe00 CUI:0 CUD:0 CURLP:(nil)
ASYNC LP:(nil) MODE:1 STATE:1 CQ LEN:0 MAX:0
PRIMARY LP: VSAN:0 CH:0x0 CHI:0 CU:0x0 CUI:0
ALTERNATE LP: VSAN:0 CH:0x0 CHI:0 CU:0x0 CUI:0

Logical Path:0x80b9fb4
VSAN:20 CH:0x200600 CHI:15 CU:0x20fe00 CUI:0 STATE:1 FLAGS:0x1
LINK: OH:0x0 OC:0x0 IH:0x0 IC:0x0
DEV: OH:0x0 OC:0x0 IH:0x0 IC:0x0
SENSE: 00 00 00 00 00 00 00 00 46
        30 20 00 00 00 00 00 00
        00 00 00 00 00 00 00 00
        00 00 00 00 00 00 00 00
IUI:0x0 DHF:0x0 CCW:0x0 TOKEN:0x0 PCCW:0x0 FCCW:0x0 PTOKEN:0x0 FTOKEN:0x0
CMD:0x0 CCW_FLAGS:0x0 CCW_COUNT:0 CMD_FLAGS:0x0 PRIO:0x0 DATA_COUNT:0
STATUS:0x0 FLAGS:0x0 PARAM:0x0 QTP:0x0 DTP:0x0
CQ LEN:0 MAX:0 DESTATUS:0x0
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying Buffer Information

In [Example 27-13](#), the `Key Counter` column displays the 32-bit value maintained by Cisco MDS switches. This value is incremented when any port changes state in that VSAN. The key counter (a 32-bit value) is incremented when a FICON-related configuration is changed. Host programs can increment this value at the start of the channel program and then perform operations on multiple ports. The director history buffer keeps a log of which port address configuration was changed for each key-counter value.

The director history buffer provides a mechanism to determine the change in the port state from the previous time when a value was contained in the key counter.

Example 27-13 Displays the History Buffer for the Specified VSAN

```
switch# show ficon vsan 20 director-history
```

```
Director History Buffer for vsan 20
```

```
-----
Key Counter          Ports Address
                    Changed
-----
74556                43
74557                44
74558                45
74559                46
74560                47
74561                48
74562                49
74563                50
74564                51
74565                52
74566                53
74567                54
74568                55
74569                56
74570                57
74571                58
74572                59
74573                60
74574                61
74575                62
74576                63
74577                64
74578
74579
74580                1-3, 5, 10, 12, 14-16, 34-40, 43-45, 47-54, 56-57, 59-64
74581                3, 5
74582                64
74583
74584                1-3, 10, 12, 14-16, 34-40, 43-45, 47-54, 56-57, 59-64
74585                1
74586                2
74587                3
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying FICON Information in the Running Configuration

Example 27-14 displays the FICON-related information in the running configuration.

Example 27-14 Displays the Running Configuration Information

```
switch# show running-config
Building Configuration ...
in-order-guarantee
vsan database
    vsan 11 name "FICON11" loadbalancing src-dst-id
    vsan 75 name "FICON75" loadbalancing src-dst-id

fcdomain domain 11 static vsan 11
fcdomain domain 119 static vsan 75

fcdroplateness network 100 vsan 11
fcdroplateness network 500 vsan 75

fabric-binding enable
fabric-binding database vsan 11
    swmn 20:00:00:0d:ec:01:20:c0 domain 10
fabric-binding database vsan 75
    swmn 20:00:00:0d:ec:00:d6:40 domain 117
fabric-binding activate vsan 11
fabric-binding activate vsan 75

ficon vsan 75

interface port-channel 1
    ficon portnumber 0x80
    switchport mode E

snmp-server user mblair network-admin auth md5 0x688fa3a2e51ba5538211606e59ac292
7 priv 0x688fa3a2e51ba5538211606e59ac2927 localizedkey
snmp-server user wwilson network-admin auth md5 0x688fa3a2e51ba5538211606e59ac29
27 priv 0x688fa3a2e51ba5538211606e59ac2927 localizedkey
snmp-server host 171.71.187.101 traps version 2c public udp-port 1163
snmp-server host 172.18.2.247 traps version 2c public udp-port 2162

vsan database
    vsan 75 interface fc1/1
...
interface mgmt0
    ip address 172.18.47.39 255.255.255.128
    switchport speed 100
    switchport duplex full

no system health

ficon vsan 75
file IPL
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying FICON Information in the Startup Configuration

[Example 27-15](#) displays the FICON-related information in the startup configuration.

Example 27-15 Displays the Startup Configuration

```
switch# show startup-config
...
ficon vsan 2
file IPL
```

[Example 27-16](#) displays the switch response to an implicitly-issued copy running start command. In this case, only a binary configuration is saved until you explicitly issue the **copy running start** command again (see [Table 27-2](#))

Example 27-16 Displays the Startup Configuration Status

```
switch# show startup-config
No ASCII config available since configuration was last saved internally
on account of 'active=saved' mode.
Please perform an explicit 'copy running startup' to get ASCII configuration
```

Displaying FICON-Related Log Information

[Example 27-17](#) and [Example 27-18](#) display the logging information for FICON-related configurations.

Example 27-17 Displays Logging Levels for the FICON Feature

```
switch# show logging level ficon
```

| Facility | Default Severity | Current Session Severity |
|----------------|------------------|--------------------------|
| ficon | 2 | 2 |
| 0(emergencies) | 1(alerts) | 2(critical) |
| 3(errors) | 4(warnings) | 5(notifications) |
| 6(information) | 7(debugging) | |

Example 27-18 Displays FICON -Related Log File Contents

```
switch# show logging logfile
...
2004 Feb 25 15:38:50 vegas6 %PORT-5-IF_UP: %$VSAN 75: 2004 Wed Feb 25 13:22:04.
131183%$ Interface fc1/8 is up in mode F
2004 Feb 25 15:38:50 vegas6 %PORT-5-IF_UP: %$VSAN 75: 2004 Wed Feb 25 13:22:04.
131217%$ Interface fc1/9 is up in mode F
...
2004 Feb 25 15:39:09 vegas6 %PORT-5-IF_TRUNK_UP: %$VSAN 75: 2004 Wed Feb 25 13:
22:23.131121%$ Interface fc2/1, vsan 75 is up
2004 Feb 25 15:39:09 vegas6 %PORT-5-IF_TRUNK_UP: %$VSAN 75: 2004 Wed Feb 25 13:
22:23.131121%$ Interface fc2/2, vsan 75 is up
2004 Feb 25 15:39:09 vegas6 %PORT-5-IF_TRUNK_UP: %$VSAN 75: 2004 Wed Feb 25 13:
...
2004 Feb 25 23:22:36 vegas6 %PORT-5-IF_UP: %$VSAN 75: 2004 Wed Feb 25 21:05:42.
99916%$ Interface fc3/6 is up in mode F
2004 Feb 25 23:22:37 vegas6 %PORT-5-IF_UP: %$VSAN 75: 2004 Wed Feb 25 21:05:43.
...
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Fabric Binding Configuration

The Cisco SAN-OS Release 1.3 fabric binding feature ensures ISLs are only enabled between specified switches in the fabric binding configuration. Fabric binding is configured on a per-VSAN basis and can only be implemented in FICON VSANs. You can still perform fabric binding configuration in a non-FICON VSAN—these configurations will only come into effect after FICON is enabled.

This feature helps prevent unauthorized switches from joining the fabric or disrupting current fabric operations. It uses the Exchange Fabric Membership Data (EFMD) protocol in FICON networks to ensure that the list of authorized switches is identical in all switches in the fabric.

This section has the following topics:

- “Port Security Versus Fabric Binding” section on page 27-37
- “Fabric Binding Enforcement” section on page 27-38
- “Fabric Binding Initiation” section on page 27-38
- “Switch WWN List Configuration” section on page 27-39
- “Fabric Binding Activation” section on page 27-39
- “Saving Fabric Binding Configurations” section on page 27-40
- “Clearing the Fabric Binding Statistics” section on page 27-41
- “Deleting the Fabric Binding Database” section on page 27-41
- “Verifying Fabric Binding Configurations” section on page 27-41

Port Security Versus Fabric Binding

Port security and fabric binding are two independent features that can be configured to complement each other (see [Table 27-3](#)).

Table 27-3 *Fabric Binding and Port Security Comparison*

| Fabric Binding | Port Security |
|--|--|
| Uses a set of sWWN and a persistent Domain ID. | Uses pWWNs/nWWNs or fWWNs/switch WWNs. |
| Binds the fabric at the switch level. | Binds devices at the interface level. |
| Authorizes only the configured sWWN stored in the fabric binding database to participate in the fabric. | Allows a preconfigured set of Fibre Channel devices to logically connect to a SAN port(s). The switchport, identified by a WWN or interface number, connects to a Fibre Channel device (a host or another switch), also identified by a WWN. By binding these two devices, you lock these two ports into a group (list). |
| Activation is required on a per VSAN basis. | Activation is required on a per VSAN basis. |
| User defines specific switches that are allowed to connect to the fabric, regardless of the physical port to which the peer switch is connected. | User specifies the specific physical port(s) to which another device can connect. |
| Does not learn logging in switches. | Learns about switches or devices if in learning mode. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Port-level checking for xE-ports

- switch login uses both port binding as well as the fabric binding feature for a given VSAN.
- Binding checks are done on the port VSAN:
 - E-port security binding check is done on port VSAN.
 - TE-port security binding check is done in each allowed VSAN.

While port security complements fabric binding, they are independent features and can be enabled or disabled separately.

Fabric Binding Enforcement

To enforce fabric binding, configure the switch world wide name (sWWN) to specify the xE port connection for each switch. Enforcement of fabric binding policies are done on every activation and when the port tries to come up. However, enforcement of fabric binding at the time of activation happens only if the VSAN is a FICON VSAN. The fabric binding feature requires all sWWNs connected to a switch and their persistent domain IDs to be part of the fabric binding active database.

To configure fabric binding in each switch in the fabric, follow these steps.

-
- | | |
|---------------|---|
| Step 1 | Enable the fabric configuration feature. |
| Step 2 | Configure a list of sWWNs and their corresponding domain IDs for devices that are allowed to access the fabric. |
| Step 3 | Activate the fabric binding database. |
| Step 4 | Save the fabric binding configuration. |
| Step 5 | Verify the fabric binding configuration. |
-

Fabric Binding Initiation

The fabric binding feature must be enabled in each switch in the fabric that participates in the fabric binding. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family. The configuration and verification commands for the fabric binding feature are only available when fabric binding is enabled on a switch. When you disable this configuration, all related configurations are automatically discarded.

To enable fabric binding on any participating switch, follow these steps:

| | Command | Purpose |
|---------------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# fabric-binding enable | Enables fabric binding on that switch. |
| | switch(config)# no fabric-binding enable | Disables (default) fabric binding on that switch. |

View the status of the fabric binding feature of an fabric binding-enabled switch by issuing the **show fabric-binding status** command.

```
switch# show fabric-binding status
VSAN 1 :Activated database
VSAN 4 :No Active database
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Switch WWN List Configuration

A user-specified fabric binding list contains a list of switch WWNs (sWWNs) within a fabric. If a sWWN attempts to join the fabric, and that sWWN is not in the list or the sWWN is using a domain ID that differs from the one specified in the allowed list, the ISL between the switch and the fabric is automatically isolated in that VSAN and the switch is denied entry into the fabric.

The persistent domain ID must be specified along with the sWWN. Domain ID authorization is required in FICON VSANs where the domains are statically configured and the end devices reject a domain ID change in all switches in the fabric.

To configure a list of sWWNs and domain IDs, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fabric-binding database vsan 5 switch(config-fabric-binding)# | Enters the fabric binding submode for the specified VSAN. |
| | switch(config)# no fabric-binding database vsan 10 | Deletes the fabric binding database for the specified VSAN. |
| Step 3 | switch(config-fabric-binding)# swwn 21:00:05:30:23:11:11:11 domain 102 | Adds the sWWN and domain ID of a switch to the configured database list. |
| Step 4 | switch(config-fabric-binding)# swwn 21:00:05:30:23:1a:11:03 domain 101 | Adds the sWWN and domain ID of another switch to the configured database list. |
| Step 5 | switch(config-fabric-binding)# no swwn 21:00:15:30:23:1a:11:03 domain 101 | Deletes the sWWN and domain ID of a switch from the configured database list. |
| Step 6 | switch(config-fabric-binding)# exit switch(config)# | Exits the fabric binding submode. |

Fabric Binding Activation

The fabric binding maintains a configuration database (config-database) and an active database. The config-database is a read-write database that collects the configurations you perform. These configurations are only enforced upon activation. This activation overwrites the active database with the contents of the config database. The active database is read-only and is the database that checks each switch that attempts to log in.

By default, the fabric binding feature is not activated. You cannot activate the switch if entries existing in the config database conflict with the current state of the fabric. For example, one of the already logged in switches may be denied login by the config database. You can choose to forcefully override these situations.



Note

After activation, any already logged in switch that violates the current active database will be logged out, and all switches that were previously denied login because of fabric binding restrictions are reinitialized.

Send documentation comments to mdsfeedback-doc@cisco.com.

To activate the fabric binding feature, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fabric-binding activate vsan 1 | Activates the fabric binding database for the specified VSAN. |
| | switch(config)# no fabric-binding activate vsan 10 | Deactivates the fabric binding database for the specified VSAN. |

Forcing Fabric Binding Activation

If the database activation is rejected due to one or more conflicts listed in the previous section, you may decide to proceed with the activation by using the **force** option.

To forcefully activate the fabric binding database, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fabric-binding activate vsan 3 force | Activates the fabric binding database for the specified VSAN forcefully—even if the configuration is not acceptable. |
| | switch(config)# no fabric-binding activate vsan 1 force | Reverts to the previously configured state or to the factory default (if no state is configured). |

Saving Fabric Binding Configurations

When you save the fabric binding configuration, the config database and the active database are both saved to the startup configuration and are available after a reboot.



Caution

You cannot deactivate or disable fabric binding in a FICON-enabled VSAN.

- Use the **fabric-binding database copy vsan** command to copy from the active database to the configuration database. If the configured database is empty, this command is not accepted.
switch# **fabric-binding database copy vsan 1**
- Use the **fabric-binding database diff active vsan** command to view the differences between the active database and the config database. This command can be used when resolving conflicts.
switch# **fabric-binding database diff active vsan 1**
- Use the **fabric-binding database diff config vsan** command to obtain information on the differences between the config database and the active database.
switch# **fabric-binding database diff config vsan 1**

Send documentation comments to mdsfeedback-doc@cisco.com.

Clearing the Fabric Binding Statistics

Use the **clear fabric-binding statistics** command to clear all existing statistics from the fabric binding database for a specified VSAN.

```
switch# clear fabric-binding statistics vsan 1
```

Deleting the Fabric Binding Database

Use the **no fabric-binding** command in configuration mode to delete the configured database for a specified VSAN.

```
switch(config)# no fabric-binding database vsan 1
```

Verifying Fabric Binding Configurations

Use the **show** commands to display all fabric binding information configured on this switch (see Examples 27-19 to 27-27).

Example 27-19 Displays Configured Fabric Binding Database Information

```
switch# show fabric-binding database
-----
Vsan    Logging-in Switch WWN      Domain-id
-----
1       21:00:05:30:23:11:11:11    0x66 (102)
1       21:00:05:30:23:1a:11:03    0x19 (25)
1       20:00:00:05:30:00:2a:1e    0xea (234)
4       21:00:05:30:23:11:11:11    0x66 (102)
4       21:00:05:30:23:1a:11:03    0x19 (25)
61      21:00:05:30:23:1a:11:03    0x19 (25)
61      21:00:05:30:23:11:11:11    0x66 (102)
[Total 7 entries]
```

Example 27-20 Displays Active Fabric Binding Information

```
switch# show fabric-binding database active
-----
Vsan    Logging-in Switch WWN      Domain-id
-----
1       21:00:05:30:23:11:11:11    0x66 (102)
1       21:00:05:30:23:1a:11:03    0x19 (25)
1       20:00:00:05:30:00:2a:1e    0xea (234)
61      21:00:05:30:23:1a:11:03    0x19 (25)
61      21:00:05:30:23:11:11:11    0x66 (102)
61      20:00:00:05:30:00:2a:1e    0xef (239)
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 27-21 Displays Active VSAN-Specific Fabric Binding Information

```
switch# show fabric-binding database active vsan 61
-----
Vsan    Logging-in Switch WWN      Domain-id
-----
61      21:00:05:30:23:1a:11:03      0x19(25)
61      21:00:05:30:23:11:11:11      0x66(102)
61      20:00:00:05:30:00:2a:1e      0xef(239)
[Total 3 entries]
```

Example 27-22 Displays Configured VSAN-Specific Fabric Binding Information

```
switch# show fabric-binding database vsan 4
-----
Vsan    Logging-in Switch WWN      Domain-id
-----
4       21:00:05:30:23:11:11:11      0x66(102)
4       21:00:05:30:23:1a:11:03      0x19(25)
[Total 2 entries]
```

Example 27-23 Displays Fabric Binding Statistics

```
switch# show fabric-binding statistics
Statistics For VSAN: 1
-----
Number of sWWN permit: 0
Number of sWWN deny  : 0

Total Logins permitted : 0
Total Logins denied   : 0
Statistics For VSAN: 4
-----
Number of sWWN permit: 0
Number of sWWN deny  : 0

Total Logins permitted : 0
Total Logins denied   : 0
Statistics For VSAN: 61
-----
Number of sWWN permit: 0
Number of sWWN deny  : 0

Total Logins permitted : 0
Total Logins denied   : 0
Statistics For VSAN: 345
-----
Number of sWWN permit: 0
Number of sWWN deny  : 0

Total Logins permitted : 0
Total Logins denied   : 0
Statistics For VSAN: 346
-----
Number of sWWN permit: 0
Number of sWWN deny  : 0

Total Logins permitted : 0
Total Logins denied   : 0
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```

Statistics For VSAN: 347
-----
Number of sWWN permit: 0
Number of sWWN deny  : 0

Total Logins permitted : 0
Total Logins denied    : 0
Statistics For VSAN: 348
-----
Number of sWWN permit: 0
Number of sWWN deny  : 0

Total Logins permitted : 0
Total Logins denied    : 0
Statistics For VSAN: 789
-----
Number of sWWN permit: 0
Number of sWWN deny  : 0

Total Logins permitted : 0
Total Logins denied    : 0
Statistics For VSAN: 790
-----
Number of sWWN permit: 0
Number of sWWN deny  : 0

Total Logins permitted : 0
Total Logins denied    : 0

```

Example 27-24 Displays Fabric Binding Status for Each VSAN

```

switch# show fabric-binding status
VSAN 1 :Activated database
VSAN 4 :No Active database
VSAN 61 :Activated database
VSAN 345 :No Active database
VSAN 346 :No Active database
VSAN 347 :No Active database
VSAN 348 :No Active database
VSAN 789 :No Active database
VSAN 790 :No Active database

```

Example 27-25 Displays Fabric Binding Violations

```

switch# show fabric-binding violations
-----
VSAN Switch WWN [domain] Last-Time [Repeat count] Reason
-----
3 20:00:00:05:30:00:4a:1e [*] Nov 25 05:44:58 2003 [2] sWWN not found
3 20:00:00:05:30:00:4a:1e [0xeb] Nov 25 05:46:14 2003 [2] Domain mismatch
4 20:00:00:05:30:00:4a:1e [*] Nov 25 05:46:25 2003 [1] Database mismatch

```



Note

In VSAN 100, the * indicates that the sWWN itself was not found in the list. In VSAN 2, the sWWN was found in the list, but has a domain ID mismatch.

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 27-26 Displays EFMD Statistics

```
switch# show fabric-binding efmd statistics

EFMD Protocol Statistics for VSAN 1
-----
Merge Requests -> Transmitted : 0 , Received : 0
Merge Accepts  -> Transmitted : 0 , Received : 0
Merge Rejects  -> Transmitted : 0 , Received : 0
Merge Busy     -> Transmitted : 0 , Received : 0
Merge Errors   -> Transmitted : 0 , Received : 0

EFMD Protocol Statistics for VSAN 4
-----
Merge Requests -> Transmitted : 0 , Received : 0
Merge Accepts  -> Transmitted : 0 , Received : 0
Merge Rejects  -> Transmitted : 0 , Received : 0
Merge Busy     -> Transmitted : 0 , Received : 0
Merge Errors   -> Transmitted : 0 , Received : 0

EFMD Protocol Statistics for VSAN 61
-----
Merge Requests -> Transmitted : 0 , Received : 0
Merge Accepts  -> Transmitted : 0 , Received : 0
Merge Rejects  -> Transmitted : 0 , Received : 0
Merge Busy     -> Transmitted : 0 , Received : 0
Merge Errors   -> Transmitted : 0 , Received : 0
```

Example 27-27 Displays EFMD Statistics for a Specified VSAN

```
switch# show fabric-binding efmd statistics vsan 4

EFMD Protocol Statistics for VSAN 4
-----
Merge Requests -> Transmitted : 0 , Received : 0
Merge Accepts  -> Transmitted : 0 , Received : 0
Merge Rejects  -> Transmitted : 0 , Received : 0
Merge Busy     -> Transmitted : 0 , Received : 0
Merge Errors   -> Transmitted : 0 , Received : 0
```

Displaying RLIR Information

The Registered Link Incident Report (RLIR) application provides a method for a switchport to send an Link Incident Record (LIR) to a registered Nx-port. It is a highly-available application.

When a LIR is detected in FICON-enabled switches in the Cisco MDS 9000 Family from a RLIR Extended Link Service (ELS). It sends that record to the members in it's Established Registration List (ERL).

In case of multi-switch topology, a Distribute Registered Link Incident Record (DRLIR) Inter Link Service (ILS) is sent to all reachable remote domains along with the RLIR ELS. On receiving the DRLIR ILS, the switch extracts the RLIR ELS and sends to the members of the ERL.

The Nx-ports interested in receiving the RLIR ELS send Link Incident Record Registration (LIRR) ELS request to the management server on the switch. The RLIRs are processed on a per-VSAN basis.

The RLIR data is written to persistent storage when the **copy running-config startup-config** command is issued.

Send documentation comments to mdsfeedback-doc@cisco.com.

The **show rlir statistics** command displays the complete statistics of LIRR, RLIR, and DRLIR frames. It lists the number of frames received, sent, and rejected. Specify the VSAN ID to obtain VSAN statistics for a specific VSAN. If you do not specify the VSAN ID, then the statistics are shown for all active VSANs (see Examples 27-28 and 27-29).

Example 27-28 Displays RLIR Statistics for All VSANs

```
switch# show rlir statistics

Statistics for VSAN: 1
-----

Number of LIRR received      = 0
Number of LIRR ACC sent      = 0
Number of LIRR RJT sent      = 0
Number of RLIR sent          = 0
Number of RLIR ACC received  = 0
Number of RLIR RJT received  = 0
Number of DRLIR received     = 0
Number of DRLIR ACC sent     = 0
Number of DRLIR RJT sent     = 0
Number of DRLIR sent         = 0
Number of DRLIR ACC received = 0
Number of DRLIR RJT received = 0

Statistics for VSAN: 100
-----

Number of LIRR received      = 26
Number of LIRR ACC sent      = 26
Number of LIRR RJT sent      = 0
Number of RLIR sent          = 815
Number of RLIR ACC received  = 815
Number of RLIR RJT received  = 0
Number of DRLIR received     = 417
Number of DRLIR ACC sent     = 417
Number of DRLIR RJT sent     = 0
Number of DRLIR sent         = 914
Number of DRLIR ACC received = 828
Number of DRLIR RJT received = 0
```

Example 27-29 Displays RLIR Statistics for a Specified VSAN

```
switch# show rlir statistics vsan 4

Statistics for VSAN: 4
-----

Number of LIRR received      = 0
Number of LIRR ACC sent      = 0
Number of LIRR RJT sent      = 0
Number of RLIR sent          = 0
Number of RLIR ACC received  = 0
Number of RLIR RJT received  = 0
Number of DRLIR received     = 0
Number of DRLIR ACC sent     = 0
Number of DRLIR RJT sent     = 0
Number of DRLIR sent         = 0
Number of DRLIR ACC received = 0
Number of DRLIR RJT received = 0
```

Send documentation comments to mdsfeedback-doc@cisco.com.

The **show rlr erl** command shows the list of Nx-ports that are registered to receive the RLIRs with the switch. If the VSAN ID is not specified, the details are shown for all active VSANs (see Examples 27-30 and 27-31).

Example 27-30 Displays All ERLs

```
switch# show rlr erl

Established Registration List for VSAN: 2
-----
FC-ID          LIRR FORMAT    REGISTERED FOR
-----
0x0b0200      0x18           always receive
Total number of entries = 1

Established Registration List for VSAN: 100
-----
FC-ID          LIRR FORMAT    REGISTERED FOR
-----
0x0b0500      0x18           conditional receive
0x0b0600      0x18           conditional receive
Total number of entries = 2
```

In Example 27-30, if the Registered For column states that an FC ID is conditional receive, the source port is registered as a valid recipient of subsequent RLIRs. This source port is selected as an RLIR recipient only if no other ERL recipient is selected.

In Example 27-30, if the Registered For column states that an FC ID is always receive, the source port is registered as a valid recipient of subsequent RLIRs. This source port is always selected as an RLIR recipient.



Note

If an *always receive* RLIR is not registered for any N-port or if the delivery of an RLIR fails for one of those ports, then the RLIR is sent to a port registered to *conditional receive* RLIRs.

Example 27-31 Displays ERLs for the Specified VSAN

```
switch# show rlr erl vsan 100
Established Registration List for VSAN: 100
-----
FC-ID          LIRR FORMAT    REGISTERED FOR
-----
0x0b0500      0x18           conditional receive
0x0b0600      0x18           conditional receive
Total number of entries = 2
```



Note

In Examples 27-32, 27-33, and 27-34, if the host time stamp (marked by the *) is available, it is printed along with the switch time stamp. If the host time stamp is not available, only the switch time stamp is printed.

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 27-32 Displays the LIR History

```
switch# show rlir history
```

```
Link incident history
```

```
-----
*Host Time Stamp
Switch Time Stamp          Port   Interface   Link Incident
-----
*Sun Nov 30 21:47:28 2003
Sun Nov 30 13:47:55 2003      2      fc1/2      Implicit Incident
*Sun Nov 30 22:00:47 2003
Sun Nov 30 14:01:14 2003      2      fc1/2      NOS Received
*Sun Nov 30 22:00:55 2003
Sun Nov 30 14:01:22 2003      2      fc1/2      Implicit Incident
*Mon Dec 1 20:14:26 2003
Mon Dec 1 12:14:53 2003      4      fc1/4      Implicit Incident
*Mon Dec 1 20:14:26 2003
Mon Dec 1 12:14:53 2003      4      fc1/4      Implicit Incident
*Thu Dec 4 04:43:32 2003
Wed Dec 3 20:43:59 2003      2      fc1/2      NOS Received
*Thu Dec 4 04:43:41 2003
Wed Dec 3 20:44:08 2003      2      fc1/2      Implicit Incident
*Thu Dec 4 04:46:53 2003
Wed Dec 3 20:47:20 2003      2      fc1/2      NOS Received
*Thu Dec 4 04:47:05 2003
Wed Dec 3 20:47:32 2003      2      fc1/2      Implicit Incident
*Thu Dec 4 04:48:07 2003
Wed Dec 3 20:48:34 2003      2      fc1/2      NOS Received
*Thu Dec 4 04:48:39 2003
Wed Dec 3 20:49:06 2003      2      fc1/2      Implicit Incident
*Thu Dec 4 05:02:20 2003
Wed Dec 3 21:02:47 2003      2      fc1/2      NOS Received
*Thu Dec 4 05:02:29 2003
Wed Dec 3 21:02:56 2003      2      fc1/2      Implicit Incident
*Thu Dec 4 05:02:47 2003
Wed Dec 3 21:03:14 2003      4      fc1/4      NOS Received
...
```

Example 27-33 Displays Recent LIRs for a Specified Interface

```
switch# show rlir recent interface fc1/1-16
```

```
Recent link incident records
```

```
-----
*Host Time Stamp
Switch Time Stamp          Port   Interface   Link Incident
-----
*Thu Dec 4 05:02:29 2003
Wed Dec 3 21:02:56 2003      2      fc1/2      Implicit Incident
*Thu Dec 4 05:02:54 2003
Wed Dec 3 21:03:21 2003      4      fc1/4      Implicit Incident
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 27-34 Displays Recent LIRs for a Specified Port Number

```
switch# show rlir recent portnumber 1-16
Recent link incident records
-----
*Host Time Stamp
Switch Time Stamp          Port   Interface   Link Incident
-----
*Thu Dec 4 05:02:29 2003
Wed Dec 3 21:02:56 2003      2       fc1/2      Implicit Incident
*Thu Dec 4 05:02:54 2003
Wed Dec 3 21:03:21 2003      4       fc1/4      Implicit Incident
```

Clearing RLIR Information

Use the **clear rlir statistics** command to clear all existing statistics for a specified VSAN.

```
switch# clear rlir statistics vsan 1
```

Use the **clear rlir history** command to clear the RLIR history where all link incident records are logged for all interfaces.

```
switch# clear rlir history
```

Use the **clear rlir recent interface** command to clear the most recent RLIR information for a specified interface.

```
switch# clear rlir recent interface fc 1/2
```

Use the **clear rlir recent portnumber** command to clear the most recent RLIT information for a specified port number.

```
switch# clear rlir recent portnumber 16
```

Default Settings

Table 27-4 lists the default settings for FICON features.

Table 27-4 Default FICON Settings

| Parameters | Default |
|-----------------------|--|
| FICON feature | Disabled. |
| Port numbers | Are the same as port addresses. |
| FC ID last byte value | 0 (zero). |
| EBCDIC format option | US-Canada. |
| Switch offline state | Hosts are allowed to move the switch to an offline state. |
| Mainframe users | Allowed to configure FICON parameters on Cisco MDS switches. |
| Clock in each VSAN | Same as the switch hardware clock. |
| Host clock control | Allows host to set the clock on this switch. |
| SNMP users | Configure FICON parameters. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 27-4 **Default FICON Settings (continued)**

| Parameters | Default |
|------------------|---|
| Port address | Not blocked |
| Prohibited ports | 90–253 and 255 for the Cisco MDS 9200 Series switches. 250–253 and 255 for the Cisco MDS 9500 Series switches. |

Table 27-5 lists the default settings for fabric binding features.

Table 27-5 **Default Fabric Binding Settings**

| Parameters | Default |
|----------------|-----------|
| Fabric binding | Disabled. |

Send documentation comments to mdsfeedback-doc@cisco.com.



CHAPTER 28

Configuring IP Storage

Cisco MDS 9000 Family IP storage (IPS) services extend the reach of Fibre Channel SANs by using open-standard, IP-based technology. The switch connects separated SAN islands using Fibre Channel over IP (FCIP), and it allows IP hosts to access Fibre Channel storage using the iSCSI protocol.



Note

FCIP and iSCSI features are specific to the IPS module and are available in Cisco MDS 9200 Switches or Cisco MDS 9500 Directors running Cisco MDS SAN-OS Release 1.1 or later.

The Cisco MDS 9216I switch and the 14/2 Multiprotocol Services (MPS-14/2) module also allow you to use Fibre Channel, FCIP, and iSCSI features. The MPS-14/2 module is available for use in any switch in the Cisco MDS 9200 Series or Cisco MDS 9500 Series running Cisco MDS SAN-OS Release 2.0(1b) or later.

This chapter includes the following sections:

- [Services Modules, page 28-2](#)
- [Supported Hardware, page 28-4](#)
- [Configuring Gigabit Ethernet Interfaces, page 28-4](#)
- [Configuring FCIP, page 28-19](#)
- [Configuring iSCSI, page 28-51](#)
- [iSCSI Authentication Setup Guidelines and Scenarios, page 28-93](#)
- [Configuring iSCSI Storage Name Services, page 28-106](#)
- [IPS Module Core Dumps, page 28-120](#)
- [Default Settings, page 28-121](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

Services Modules


The IP Storage services module (IPS module) and the MPS-14/2 module allow you to use FCIP and iSCSI features. Both modules integrate seamlessly into the Cisco MDS 9000 Family, and support the full range of features available on other switching modules, including VSANs, security, and traffic management. The following types storage services modules are currently available for use in any switch in the Cisco MDS 9200 Series or in the Cisco MDS 9500 Series:

- The 4-port, hot-swappable IPS module (IPS-4) has four Gigabit Ethernet ports.
- The 8-port, hot-swappable IPS module (IPS-8) has eight Gigabit Ethernet ports.
- The MPS-14/2 module has 14 Fibre Channel ports (numbered 1 through 14) and two Gigabit Ethernet ports (numbered 1 and 2)

Gigabit Ethernet ports in these modules can be configured to support FCIP protocol, iSCSI protocol, or both protocols simultaneously.

- FCIP—FCIP transports Fibre Channel frames transparently over an IP network between two Cisco MDS 9000 Family switches or other FCIP standards-compliant devices. [Figure 28-1](#) shows how the IPS module is used in different FCIP scenarios.


Figure 28-1 FCIP Scenarios



- iSCSI—The IPS module provides IP hosts access to Fibre Channel storage devices. The IP host sends SCSI commands encapsulated in iSCSI protocol data units (PDUs) to a Cisco MDS 9000 Family switch IPS port over a TCP/IP connection. At this point, the commands are routed from an IP network into a Fibre Channel network and forwarded to the intended target. [Figure 28-2](#) depicts the iSCSI scenarios in which the IPS module is used.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 28-2 iSCSI Scenarios



Module Status Verification

After inserting the module, verify the status of the module using the **show module** command:

```
switch# show module
```

| Mod | Ports | Module-Type | Model | Status |
|-----|-------|---------------------------------|------------------|--------------------------|
| 1 | 0 | Caching Services Module | DS-X9560-SMAP | ok |
| 2 | 8 | IP Storage Services Module | DS-X9308-SMIP | ok <-----IPS-8 module |
| 4 | 16 | 2x1GE IPS, 14x1/2Gbps FC Module | DS-X9216i-K9-SUP | ok <-----MPS-14/2 module |
| 5 | 0 | Supervisor/Fabric-1 | DS-X9530-SF1-K9 | active * |
| 6 | 0 | Supervisor/Fabric-1 | DS-X9530-SF1-K9 | ha-standby |
| 9 | 4 | IP Storage Services Module | DS-X9304-SMIP | ok <-----IPS-4 module |

| Mod | Sw | Hw | World-Wide-Name(s) (WWN) |
|-----|--------|-------|--|
| 1 | 2.0(1) | 0.201 | 20:41:00:0b:fd:44:68:c0 to 20:48:00:0b:fd:44:68:c0 |
| 2 | 2.0(1) | 0.201 | 20:41:00:0b:fd:44:68:c0 to 20:48:00:0b:fd:44:68:c0 |
| 4 | 2.0(1) | 0.201 | 20:c1:00:05:30:00:07:1e to 20:d0:00:05:30:00:07:1e |
| 5 | 2.0(1) | 0.0 | -- |
| 6 | 2.0(1) | 0.0 | -- |
| 9 | 2.0(1) | 0.1 | 22:01:00:05:30:00:07:1e to 22:04:00:05:30:00:07:1e |

| Mod | Application Image Description | Application Image Version |
|-----|-------------------------------|---------------------------|
| 1 | svc-node1 | 1.3 (5M) |
| 1 | svc-node2 | 1.3 (5M) |

| Mod | MAC-Address(es) | Serial-Num |
|-----|--|-------------|
| 1 | 00-05-30-01-49-c2 to 00-05-30-01-4a-46 | JAB073907EP |
| 2 | 00-05-30-00-9d-d2 to 00-05-30-00-9d-de | JAB064605a2 |
| 4 | 00-05-30-01-7f-32 to 00-05-30-01-7f-38 | JAB081405AM |
| 5 | 00-05-30-00-2c-4e to 00-05-30-00-2c-52 | JAB06350B1M |
| 6 | 00-05-30-00-19-66 to 00-05-30-00-19-6a | JAB073705GL |
| 9 | 00-0d-bc-2f-d6-00 to 00-0d-bc-2f-d6-08 | JAB080804TN |

* this terminal session

Send documentation comments to mdsfeedback-doc@cisco.com.

IPS Module Upgrade



Caution

A software upgrade is only disruptive for the IPS module. The SAN-OS software continues to support nondisruptive software upgrades for Fibre Channel modules in the switch and for the switch itself.

IPS modules use a rolling upgrade install mechanism where each module in a given switch can only be upgraded in sequence. To guarantee a stable state, each IPS module in a switch requires a 5-minute delay before the next IPS module is upgraded.

MPS-14/2 Module Upgrade



Caution

A software upgrade is only partially disruptive for the MPS-14/2 module. The SAN-OS software continues to support nondisruptive software upgrades for Fibre Channel modules in the switch and for the switch itself.

The MPS-14/2 modules have 14 Fibre Channel ports (nondisruptive upgrade) and 2 Gigabit Ethernet ports (disruptive upgrade). MPS-14/2 modules use a rolling upgrade install mechanism for the two Gigabit Ethernet ports where each module in a given switch can only be upgraded in sequence. To guarantee a stable state, each MPS-14/2 module in a switch requires a 5-minute delay before the next module is upgraded.

Supported Hardware

You can configure the FCIP and iSCSI features using one of more of the following hardware:

- IPS-4 and IPS-8 modules (refer to the *Cisco MDS 9200 Series Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide* for more information)
- MPS-14/2 module (refer to the *Cisco MDS 9200 Series Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide* for more information).



Note

In both the MPS-14/2 module and the Cisco MDS 9216i integrated supervisor module, the port numbering differs for the Fibre Channel and the Gigabit Ethernet ports. The Fibre Channel ports are numbered from 1 through 14 and the Gigabit Ethernet ports are numbered as 1 and 2.

- Cisco MDS 9216i Switch (refer to the *Cisco MDS 9200 Series Hardware Installation Guide*).

Configuring Gigabit Ethernet Interfaces

Both FCIP and iSCSI rely on TCP/IP for network connectivity. On each IPS module or MPS-14/2 module, connectivity is provided in the form of Gigabit Ethernet interfaces that are appropriately configured. This section covers the steps required to configure IP for subsequent use by FCIP and iSCSI.

Send documentation comments to mdsfeedback-doc@cisco.com.

A new port mode, called IPS, is defined for Gigabit Ethernet ports on each IPS module or MPS-14/2 module. IP storage ports are implicitly set to IPS mode, so it can only be used to perform iSCSI and FCIP storage functions. IP storage ports do not bridge Ethernet frames or route other IP packets.

Each IPS port represents a single virtual Fibre Channel host in the Fibre Channel SAN. All the iSCSI hosts connected to this IPS port are merged and multiplexed via the single Fibre Channel host.

In large scale iSCSI deployments where the Fibre Channel storage subsystems require explicit LUN access control for every host device, use of proxy-initiator mode simplifies the configuration.




Tip

Gigabit Ethernet ports on any IPS module or MPS-14/2 module should not be configured in the same Ethernet broadcast domain as the management Ethernet port—they should be configured in a different broadcast domain, either by using separate standalone hubs or switches or by using separate VLANs.

Configuring a Basic Gigabit Ethernet Interface

Figure 28-3 shows an example of a basic Gigabit Ethernet configuration.

Figure 28-3 Gigabit Ethernet Configuration



To configure the Gigabit Ethernet interface for the scenario in Figure 28-3, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# interface gigabitethernet 2/2 switch(config-if)# | Enters the interface configuration mode on the Gigabit Ethernet interface (slot 2, port 2). |
| Step 3 | switch(config-if)# ip address 10.1.1.100 255.255.255.0 | Enters the IP address (10.1.1.100) and subnet mask (255.255.255.0) for the Gigabit Ethernet interface. |
| Step 4 | switch(config-if)# no shutdown | Enables the interface. |

Configuring Interface Descriptions

See the “[Interface Descriptions](#)” section on page 12-12 for details on configuring the switchport description for any interface.

Configuring Beacon Mode

See the “[Beacon Mode](#)” section on page 12-16 for details on configuring the beacon mode for any interface.

Send documentation comments to mdsfeedback-doc@cisco.com.

Configuring Auto-Negotiation

By default, auto-negotiation is enabled all Gigabit Ethernet interface. You can enable or disable auto-negotiation for a specified Gigabit Ethernet interface. When auto-negotiation is enabled, the port automatically detects the speed or pause method, and duplex of incoming signals based on the link partner. You can also detect link up conditions using the auto-negotiation feature.

To configure auto-negotiation, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# interface gigabitethernet 2/2 switch(config-if)# | Enters the interface configuration mode on the Gigabit Ethernet interface (slot 2, port 2). |
| Step 3 | switch(config-if)# switchport auto-negotiate | Enables auto-negotiation for this Gigabit Ethernet interface (default). |
| | switch(config-if)# no switchport auto-negotiate | Disables auto-negotiation for this Gigabit Ethernet interface. |

Configuring the MTU Frame Size

You can configure the interfaces on a switch to transfer large (or jumbo) frames on a port. The default IP maximum transmission unit (MTU) frame size is 1500 bytes for all Ethernet ports. By configuring jumbo frames on a port, the MTU size can be increased up to 9000 bytes.



Note

The minimum MTU size is 576 bytes.



Tip

MTU changes are disruptive, all FCIP links and iSCSI sessions flap when the software detects a change in the MTU size.

You do not need to explicitly issue the **shutdown** and **no shutdown** commands.

To configure the MTU frame size, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# interface gigabitethernet 2/2 switch(config-if)# | Enters the interface configuration mode on the Gigabit Ethernet interface (slot 2, port 2). |
| Step 3 | switch(config-if)# switchport mtu 3000 | Changes the MTU size to 3000 bytes. The default is 1500 bytes. |

Configuring Promiscuous Mode

You can enable or disable promiscuous mode on a specific Gigabit Ethernet interface. By enabling the promiscuous mode, the Gigabit Ethernet interface receives all the packets and the software then filters and discards the packets that are not destined for that Gigabit Ethernet interface.

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure the promiscuous mode, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# interface gigabitethernet 2/2 switch(config-if)# | Enters the interface configuration mode on the Gigabit Ethernet interface (slot 2, port 2). |
| Step 3 | switch(config-if)# switchport promiscuous-mode on | Enables promiscuous mode for this Gigabit Ethernet interface. The default is off . |
| | switch(config-if)# switchport promiscuous-mode off | Disables (default) promiscuous mode for this Gigabit Ethernet interface. |
| | switch(config-if)# no switchport promiscuous-mode | Disables (default) the promiscuous mode for this Gigabit Ethernet interface. |

About VLANs for Gigabit Ethernet

Virtual LANs (VLANs) create multiple virtual Layer 2 networks over a physical LAN network. VLANs provide traffic isolation, security, and broadcast control.

Gigabit Ethernet ports automatically recognize Ethernet frames with IEEE 802.1Q VLAN encapsulation. If you need to have traffic from multiple VLANs terminated on one Gigabit Ethernet port, configure subinterfaces—one for each VLAN.



Note

If the IPS module or MPS-14/2 module is connected to a Cisco Ethernet switch, and you need to have traffic from multiple VLANs coming to one IPS port, verify the following requirements on the Ethernet switch:

- The Ethernet switch port connected to the IPS module or MPS-14/2 module is configured as a trunking port.
- The encapsulation is set to 802.1Q and not ISL, which is the default.

Use the VLAN ID as a subscription to the Gigabit Ethernet interface name to create the subinterface name (the <slot-number>/<port-number>.<VLAN-ID>).

Configuring the VLAN Subinterface

To configure a VLAN subinterface (VLAN ID), follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# interface gigabitethernet 2/2.100 switch(config-if)# | Specifies the subinterface on which 802.1Q is used (slot 2, port 2, VLAN ID 100). Note The subinterface number, 100 in this example, is the VLAN ID. The VLAN ID ranges from 1 to 4093. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|---|--|
| Step 3 | <code>switch(config-if)# ip address 10.1.1.101 255.255.255.0</code> | Enters the IP address (10.1.1.100) and IP mask (255.255.255.0) for the Gigabit Ethernet interface. |
| Step 4 | <code>switch(config-if)# no shutdown</code> | Enables the interface. |

Interface Subnet Requirements

Gigabit Ethernet interfaces (major), subinterfaces (VLAN ID), and management interfaces (mgmt 0) can be configured in the same or different subnet depending on the configuration (see [Table 28-1](#)).

Table 28-1 Subnet Requirements for Interfaces

| Interface 1 | Interface 2 | Same Subnet Allowed | Notes |
|--------------------------|--------------------------|---------------------|--|
| Gigabit Ethernet 1/1 | Gigabit Ethernet 1/2 | Yes | Two major interfaces can be configured in the same or different subnets. |
| Gigabit Ethernet 1/1.100 | Gigabit Ethernet 1/2.100 | Yes | Two subinterfaces with the same VLAN ID can be configured in the same or different subnets. |
| Gigabit Ethernet 1/1.100 | Gigabit Ethernet 1/2.200 | No | Two subinterfaces with different VLAN IDs cannot be configured in the same subnet. |
| Gigabit Ethernet 1/1 | Gigabit Ethernet 1/1.100 | No | A subinterface cannot be configured on the same subnet as the major interface. |
| mgmt0 | Gigabit Ethernet 1/1.100 | No | The mgmt0 interface cannot be configured in the same subnet as the Gigabit Ethernet interfaces or subinterfaces. |
| mgmt0 | Gigabit Ethernet 1/1 | No | |



Note

The configuration requirements in [Table 28-1](#) also apply to Ethernet PortChannels.

Configuring Static IP Routing

To configure static IP routing (see [Figure 28-3](#)) through the Gigabit Ethernet interface, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | <code>switch# config terminal</code> <code>switch(config)#</code> | Enters configuration mode. |
| Step 2 | <code>switch(config)# ip route 10.100.1.0 255.255.255.0 10.1.1.1</code> <code>switch(config-if)#</code> | Enters the IP subnet (10.100.1.0 255.255.255.0) of the IP host and configures the next hop 10.1.1.1, which is the IP address of the router connected to the Gigabit Ethernet interface. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying the IP Route Table

The **show ips ip route interface** command takes the Gigabit Ethernet interface as a parameter and returns the route table for the interface. See [Example 28-1](#).

Example 28-1 Displays the IP Route Table

```
switch# show ips ip route interface gig 8/1
Codes: C - connected, S - static
No default gateway
C 10.1.3.0/24 is directly connected, GigabitEthernet8/1
```

Connected (C) identifies the subnet in which the interface is configured (directly connected to the interface). Static (S) identifies the static routes that go through the router.

Verifying Gigabit Ethernet Connectivity

Once the Gigabit Ethernet interfaces are connected with valid IP addresses, verify the interface connectivity on each switch. Ping the IP host using the IP address of the host to verify that the static IP route is configured correctly.



Note

If the connection fails, verify the following, and ping the IP host again:

- The IP address for the destination (IP host) is correctly configured.
- The host is active (powered on).
- The IP route is configured correctly.
- The IP host has a route to get to the Gigabit Ethernet interface subnet.
- The Gigabit Ethernet interface is in the up state.

Use the **ping** command to verify the Gigabit Ethernet connectivity (see [Example 28-2](#)). The **ping** command sends echo request packets out to a remote device at an IP address that you specify (see the [“Using the ping Command”](#) section on page 2-14).

Use the **show interface gigabitethernet** command to verify if the Gigabit Ethernet interface is up.

Example 28-2 Verifying Gigabit Ethernet Connectivity

```
switch# ping 10.100.1.25
PING 10.100.1.25 (10.100.1.25): 56 data bytes
64 bytes from 10.100.1.25: icmp_seq=0 ttl=255 time=0.1 ms
64 bytes from 10.100.1.25: icmp_seq=1 ttl=255 time=0.1 ms
64 bytes from 10.100.1.25: icmp_seq=2 ttl=255 time=0.1 ms
--- 10.100.1.25 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0.1/0.1/0.1 ms
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Gigabit Ethernet IP-ACL Guidelines



Tip

If IP-ACLs are already configured in a Gigabit Ethernet interface, you cannot add this interface to a Ethernet PortChannel group. See the “[IP Access Control Lists](#)” section on page 26-5 for information on configuring IP-ACLs.

Follow these guidelines when configuring IP-ACLs for Gigabit Ethernet interfaces:

- Only use Transmission Control Protocol (TCP, keyword **tcp**) or Internet Control Message Protocol (ICMP, keyword **icmp**).



Note

Other protocols like, User Datagram Protocol (UDP) and HTTP, are not supported in Gigabit Ethernet interfaces. Applying an ACL that contains rules for these protocols to a Gigabit Ethernet interface is allowed but those rules have no effect.

- Apply IP-ACLs to the interface before you enable an interface. This ensures that the filters are in place before traffic starts flowing.
- Be aware of the following conditions:
 - If you use the **log-deny** option, a maximum of 50 messages are logged per second.
 - The **established**, **precedence**, and **fragments** options are ignored when you apply IP-ACLs (containing these options) to Gigabit Ethernet interfaces.
 - If an IP-ACL rule applies to a pre-existing TCP connection, that rule is ignored. For example if there is an existing TCP connection between A and B and an IP-ACL which specifies dropping all packets whose source is A and destination is B is subsequently applied, it will have no effect.

Applying IP-ACLs on Gigabit Ethernet Interfaces

To apply an IP-ACL on an Gigabit Ethernet interface, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface gigabitethernet 3/1 switch(config-if)# | Configures a Gigabit Ethernet interface (3/1). |
| Step 3 | switch(config-if)# ip access-group SampleName | Applies the IP-ACL SampleName on Gigabit Ethernet 3/1 for both ingress and egress traffic (if the association does not exist already). |
| Step 4 | switch(config-if)# ip access-group SampleName1 in | Applies the IP-ACL SampleName on Gigabit Ethernet 3/1 for ingress traffic. |
| | switch(config-if)# ip access-group SampleName2 out | Applies the IP-ACL SampleName on Gigabit Ethernet 3/1 for egress traffic (if the association does not exist already). |

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying ARP Caches

You can display the ARP cache on Gigabit Ethernet interfaces.



Note

Use the physical interface, not the subinterface, for all ARP cache commands.

Use the **show ips arp interface gigabitethernet** command to display the ARP cache on the Gigabit Ethernet interfaces. This command takes the Ethernet interface as a parameter and returns the ARP cache for that interface. See [Example 28-3](#).

Example 28-3 Displays ARP Caches

```
switch# show ips arp interface gigabitethernet 7/1
Protocol      Address      Age (min)    Hardware Addr  Type   Interface
Internet      20.1.1.5     3            0005.3000.9db6 ARPA   GigabitEthernet7/1
Internet      20.1.1.10    7            0004.76eb.2ff5 ARPA   GigabitEthernet7/1
Internet      20.1.1.11    16           0003.47ad.21c4 ARPA   GigabitEthernet7/1
Internet      20.1.1.12    6            0003.4723.c4a6 ARPA   GigabitEthernet7/1
Internet      20.1.1.13    13           0004.76f0.ef81 ARPA   GigabitEthernet7/1
Internet      20.1.1.14    0            0004.76e0.2f68 ARPA   GigabitEthernet7/1
Internet      20.1.1.15    6            0003.47b2.494b ARPA   GigabitEthernet7/1
Internet      20.1.1.17    2            0003.479a.b7a3 ARPA   GigabitEthernet7/1
...
```

Clearing ARP Caches

The ARP cache can be cleared in two ways: clearing just one entry or clearing all entries in the ARP cache.

Use the **clear ips arp** command to clear the ARP cache. See [Example 28-4](#) and [Example 28-5](#).

Example 28-4 Clearing One ARP Cache Entry

```
switch# clear ips arp address 10.2.2.2 interface gigabitethernet 8/7
arp clear successful
```

Example 28-5 Clearing All ARP Cache Entries

```
switch# clear ips arp interface gigabitethernet 8/7
arp clear successful
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying Statistics

This section provides examples to verify Gigabit Ethernet and TCP/IP statistics on the IP storage ports.

Displaying Gigabit Ethernet Interface Statistics

Use the **show interface Gigabit Ethernet** command on each switch to verify that the interfaces are up and functioning as desired. See [Example 28-6](#).

Example 28-6 Displays the Gigabit Ethernet Interface

```
switch# show interface gigabitethernet 8/1
GigabitEthernet8/1 is up <-----The interface is in the up state.
  Hardware is GigabitEthernet, address is 0005.3000.a98e
  Internet address is 10.1.3.1/24
  MTU 1500 bytes, BW 1000000 Kbit
  Port mode is IPS
  Speed is 1 Gbps
  Beacon is turned off
  5 minutes input rate 744 bits/sec, 93 bytes/sec, 1 frames/sec
  5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  3343 packets input, 406582 bytes
    0 multicast frames, 0 compressed
    0 input errors, 0 frame, 0 overrun 0 fifo
  8 packets output, 336 bytes, 0 underruns
    0 output errors, 0 collisions, 0 fifo
    0 carrier errors
```

Example 28-7 Displays the Gigabit Ethernet Subinterface

```
switch# show interface gigabitethernet 4/2.100
GigabitEthernet4/2.100 is up
  Hardware is GigabitEthernet, address is 0005.3000.abcb
  Internet address is 10.1.2.100/24
  MTU 1500 bytes
  5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  0 packets input, 0 bytes
    0 multicast frames, 0 compressed
    0 input errors, 0 frame, 0 overrun 0 fifo
  1 packets output, 46 bytes, 0 underruns
    0 output errors, 0 collisions, 0 fifo
    0 carrier errors
```

Displaying Ethernet MAC Statistics

The **show ips stats mac interface gigabitethernet** command takes the main Gigabit Ethernet interface as a parameter and returns Ethernet statistics for that interface. See [Example 28-8](#).



Note

Use the physical interface, not the subinterface, to display Ethernet MAC statistics.

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 28-8 Displays Ethernet MAC Statistics

```
switch# show ips stats mac interface gigabitethernet 8/1
Ethernet MAC statistics for port GigabitEthernet8/1
  Hardware Transmit Counters
    237 frame 43564 bytes
    0 collisions, 0 late collisions, 0 excess collisions
    0 bad frames, 0 FCS error, 0 abort, 0 runt, 0 oversize
  Hardware Receive Counters
    427916 bytes, 3464 frames, 0 multicasts, 3275 broadcasts
    0 bad, 0 runt, 0 CRC error, 0 length error
    0 code error, 0 align error, 0 oversize error
  Software Counters
    3429 received frames, 237 transmit frames
    0 frames soft queued, 0 current queue, 0 max queue
    0 dropped, 0 low memory
```

Displaying DMA-Bridge Statistics

You can display direct memory access (DMA) device statistics using the **show ips stats dma-bridge interface gigabitethernet** command. This command takes the main Gigabit Ethernet interface as a parameter and returns DMA bridge statistics for that interface. See [Example 28-9](#).



Note

Use the physical interface, not the subinterface, to display DMA-bridge statistics.

Example 28-9 Displays DMA-Bridge Statistics

```
switch# show ips stats dma-bridge interface gigabitethernet 7/1
Dma-bridge ASIC Statistics for port GigabitEthernet7/1
  Hardware Egress Counters
    231117 Good, 0 bad protocol, 0 bad header cksum, 0 bad FC CRC
  Hardware Ingress Counters
    218255 Good, 0 protocol error, 0 header checksum error
    0 FC CRC error, 0 iSCSI CRC error, 0 parity error
  Software Egress Counters
    231117 good frames, 0 bad header cksum, 0 bad FIFO SOP
    0 parity error, 0 FC CRC error, 0 timestamp expired error
    0 unregistered port index, 0 unknown internal type
    0 RDL ok, 0 RDL drop (too big), 0 RDL ttl_1
    3656368645 idle poll count, 0 loopback, 0 FCC PQ, 0 FCC EQ
    Flow Control: 0 [0], 0 [1], 0 [2], 0 [3]
  Software Ingress Counters
    218255 Good frames, 0 header cksum error, 0 FC CRC error
    0 iSCSI CRC error, 0 descriptor SOP error, 0 parity error
    0 frames soft queued, 0 current Q, 0 max Q, 0 low memory
    0 out of memory drop, 0 queue full drop
    0 RDL ok, 0 RDL drop (too big)
    Flow Control: 0 [0], 0 [1], 0 [2], 0 [3]
```

This output shows all Fibre Channel frames that ingress or egress from the Gigabit Ethernet port.

Displaying TCP/IP Statistics

Use the **show ips stats ip interface gigabitethernet** to display and verify IP statistics. This command takes the main Ethernet interface as a parameter and returns the IP statistics for that interface. See [Example 28-10](#).

Send documentation comments to mdsfeedback-doc@cisco.com.



Note

Use the physical interface, not the subinterface, to display TCP/IP statistics.

Example 28-10 Displays IP Statistics

```
switch# show ips stats ip interface gigabitethernet 4/1
Internet Protocol Statistics for port GigabitEthernet4/1
  168 total received, 168 good, 0 error
  0 reassembly required, 0 reassembled ok, 0 dropped after timeout
  371 packets sent, 0 outgoing dropped, 0 dropped no route
  0 fragments created, 0 cannot fragment
```

Use the **show ips stats tcp interface gigabitethernet** to display and verify TCP statistics. This command takes the main Ethernet interface as a parameter, and shows TCP stats along with the connection list and TCP state. The **detail** option shows all information maintained by the interface. See [Example 28-11](#) and [Example 28-12](#).

Example 28-11 Displays TCP Statistics

```
switch# show ips stats tcp interface gigabitethernet 4/1
TCP Statistics for port GigabitEthernet4/1
  Connection Stats
    0 active openings, 3 accepts
    0 failed attempts, 12 reset received, 3 established
  Segment stats
    163 received, 355 sent, 0 retransmitted
    0 bad segments received, 0 reset sent
  TCP Active Connections
    Local Address      Remote Address      State      Send-Q    Recv-Q
    0.0.0.0:3260       0.0.0.0:0          LISTEN     0          0
```

Example 28-12 Displays Detailed TCP Statistics

```
switch# show ips stats tcp interface gigabitethernet 4/1 detail
TCP Statistics for port GigabitEthernet4/1
  TCP send stats
    355 segments, 37760 bytes
    222 data, 130 ack only packets
    3 control (SYN/FIN/RST), 0 probes, 0 window updates
    0 segments retransmitted, 0 bytes
    0 retransmitted while on ethernet send queue, 0 packets split
    0 delayed acks sent
  TCP receive stats
    163 segments, 114 data packets in sequence, 6512 bytes in sequence
    0 predicted ack, 10 predicted data
    0 bad checksum, 0 multi/broadcast, 0 bad offset
    0 no memory drops, 0 short segments
    0 duplicate bytes, 0 duplicate packets
    0 partial duplicate bytes, 0 partial duplicate packets
    0 out-of-order bytes, 1 out-of-order packets
    0 packet after window, 0 bytes after window
    0 packets after close
    121 acks, 37764 ack bytes, 0 ack toomuch, 4 duplicate acks
    0 ack packets left of snd_una, 0 non-4 byte aligned packets
    8 window updates, 0 window probe
    30 pcb hash miss, 0 no port, 0 bad SYN, 0 paws drops
  TCP Connection Stats
    0 attempts, 3 accepts, 3 established
    3 closed, 2 drops, 0 conn drops
    0 drop in retransmit timeout, 1 drop in keepalive timeout
```


Send documentation comments to mdsfeedback-doc@cisco.com.

```

    0 drop in persist drops, 0 connections drained
TCP Miscellaneous Stats
    115 segments timed, 121 rtt updated
    0 retransmit timeout, 0 persist timeout
    12 keepalive timeout, 11 keepalive probes
TCP SACK Stats
    0 recovery episodes, 0 data packets, 0 data bytes
    0 data packets retransmitted, 0 data bytes retransmitted
    0 connections closed, 0 retransmit timeouts
TCP SYN Cache Stats
    15 entries, 3 connections completed, 0 entries timed out
    0 dropped due to overflow, 12 dropped due to RST
    0 dropped due to ICMP unreachable, 0 dropped due to bucket overflow
    0 abort due to no memory, 0 duplicate SYN, 0 no-route SYN drop
    0 hash collisions, 0 retransmitted
TCP Active Connections
  Local Address      Remote Address      State      Send-Q      Recv-Q
  0.0.0.0:3260       0.0.0.0:0          LISTEN     0           0

```

Use the **show ips stats icmp interface gigabitethernet** to display and verify IP statistics. This command takes the main Ethernet interface as a parameter and returns the ICMP statistics for that interface. See [Example 28-13](#).

Example 28-13 Displays ICMP Statistics

```

switch# show ips stats icmp interface gigabitethernet 2/1
ICMP Statistics for port GigabitEthernet2/1
  0 ICMP messages received
  0 ICMP messages dropped due to errors
ICMP input histogram
  0 destination unreachable
  0 time exceeded
  0 parameter problem
  0 source quench
  0 redirect
  0 echo request
  0 echo reply
  0 timestamp request
  0 timestamp reply
  0 address mask request
  0 address mask reply
ICMP output histogram
  0 destination unreachable
  0 time exceeded
  0 parameter problem
  0 source quench
  0 redirect
  0 echo request
  0 echo reply
  0 timestamp request
  0 timestamp reply
  0 address mask request
  0 address mask reply

```

Configuring Gigabit Ethernet High Availability

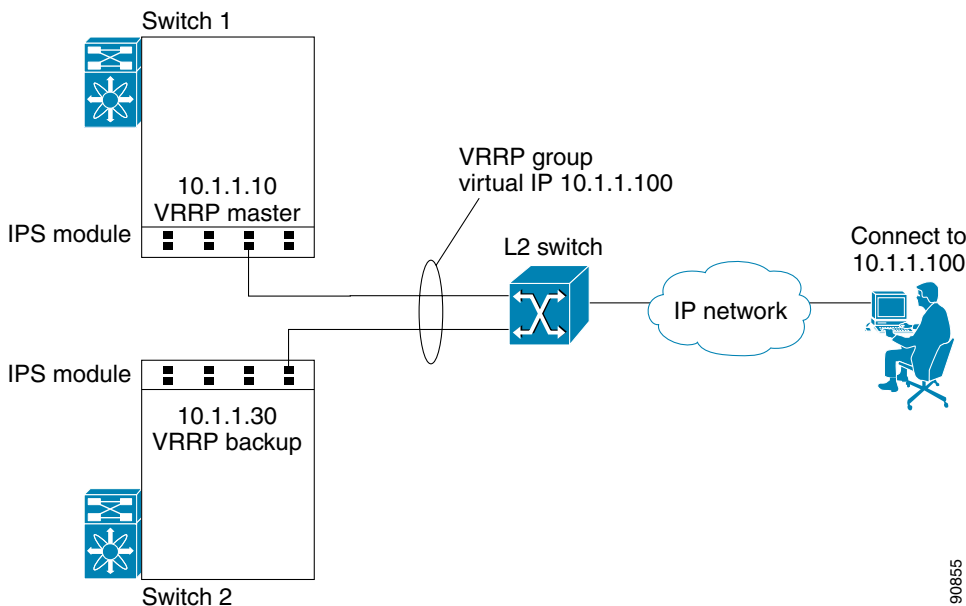
Virtual Router Redundancy Protocol (VRRP) and Ethernet PortChannels are two Gigabit Ethernet features that provide high availability for iSCSI and FCIP services.

Send documentation comments to mdsfeedback-doc@cisco.com.

VRRP for iSCSI and FCIP Services

VRRP provides a redundant alternate path to the Gigabit Ethernet port for iSCSI and FCIP services. VRRP provides IP address fail over protection to an alternate Gigabit Ethernet interface so the IP address is always available (see [Figure 28-4](#)).

Figure 28-4 VRRP Scenario



In [Figure 28-4](#), all members of the VRRP group must be IP storage Gigabit Ethernet ports. VRRP group members can be one or more of the following interfaces:

- One or more interfaces in the same IPS module or MPS-14/2 module
- Interfaces across IPS modules or MPS-14/2 modules in one switch
- Interfaces across IPS modules or MPS-14/2 modules in different switches
- Gigabit Ethernet subinterfaces
- Ethernet PortChannels and PortChannel subinterfaces

See the “[The Virtual Router Redundancy Protocol](#)” section on page 26-19.

Configuring VRRP for Gigabit Ethernet Interfaces

To configure VRRP for Gigabit Ethernet interfaces, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | <code>switch1# config terminal</code> <code>switch1(config)#</code> | Enters configuration mode. |
| Step 2 | <code>switch(config)# interface</code> <code>gigabitethernet 2/2</code> <code>switch(config-if)#</code> | Enters the interface configuration mode on the Gigabit Ethernet interface (slot 2, port 2). |
| Step 3 | <code>switch(config-if)# ip address</code> <code>10.1.1.10 255.255.255.0</code> | Enters the IP address (10.1.1.10) and IP mask (255.255.255.0) for the Gigabit Ethernet interface. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|---|---|
| Step 4 | <code>switch(config-if)# no shutdown</code> | Enables the selected interface. |
| Step 5 | <code>switch(config-if)# vrrp 100</code> <code>switch(config-if-vrrp)</code> | Creates a VR ID 100. |
| Step 6 | <code>switch(config-if-vrrp)# address 10.1.1.100</code> | Configures the virtual IP address (10.1.1.100) for the selected VRRP group (identified by the VR ID). Note The virtual IP address must be in the same subnet as the IP address of the Gigabit Ethernet interface. All members of the VRRP group must configure the same virtual IP address. |
| Step 7 | <code>switch(config-if-vrrp)# priority 10</code> | Configures the priority for the selected interface within this VRRP group. Note The interface with the highest priority is selected as the master. |
| Step 8 | <code>switch(config-if-vrrp)# no shutdown</code> | Enables the VRRP protocol on the selected interface. |

**Note**

The VRRP **preempt** option is not supported on IPS Gigabit Ethernet interfaces. However, if the virtual IP address is also the IP address for the interface, then preemption is implicitly applied.

About Ethernet PortChannel Aggregation

Ethernet PortChannels refer to the aggregation of multiple physical Gigabit Ethernet interfaces into one logical Ethernet interface to provide link redundancy and, in some cases, higher aggregated bandwidth and load balancing.

An Ethernet switch connecting to the MDS switch Gigabit Ethernet port can implement load balancing based on the IP address, IP address and UDP/TCP port number, or MAC address. Due to the load balancing scheme, the data traffic from one TCP connection is always sent out on the same physical Gigabit Ethernet port of an Ethernet PortChannel. For the traffic coming to the MDS, an ethernet switch can implement load balancing based on its IP address, its source-destination MAC address, or its IP address and port. The data traffic from one TCP connection always travels on the same physical links. To make use of both ports for the outgoing direction, multiple TCP connections are required.

All FCIP data traffic for one FCIP link is carried on one TCP connection. Consequently, the aggregated bandwidth is 1 Gbps for that FCIP link.

**Note**

The Cisco Ethernet switch's PortChannel should be configured as a static PortChannel, and not the default 802.3ad protocol.


Ethernet PortChannels can only aggregate two physical interfaces that are adjacent to each other on a given IPS module (see [Figure 28-5](#)).

**Note**

PortChannel members must be one of these combinations: ports 1–2, ports 3–4, ports 5–6, or ports 7–8.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 28-5 Ethernet PortChannel Scenario



In [Figure 28-5](#), Gigabit Ethernet ports 3 and 4 in slot 9 are aggregated into an Ethernet PortChannel. Ethernet PortChannels are not supported on MPS-14/2 modules and 9216i IPS modules.



Note

PortChannel interfaces provide configuration options for both Gigabit Ethernet and Fibre Channel. However, based on the PortChannel membership, only Gigabit Ethernet parameters or Fibre Channel parameters are applicable.

Configuring Ethernet PortChannels

The PortChannel configuration specified in [Chapter 14, “Configuring PortChannels”](#) also applies to Ethernet PortChannel configurations.

To configure Ethernet PortChannels, follow these steps:

| | Command | Purpose |
|---------------|---|--|
| Step 1 | switch1# config terminal switch1(config)# | Enters configuration mode. |
| Step 2 | switch(config)# interface port-channel 10 switch(config-if)# | Configures the specified PortChannel (10). |
| Step 3 | switch(config-if)# ip address 10.1.1.1 255.255.255.0 | Enters the IP address (10.1.1.1) and IP mask (255.255.255.0) for the PortChannel. Note A PortChannel does not have any members when first created. |
| Step 4 | switch(config-if)# no shutdown | Enables the interface. |
| Step 5 | switch(config)# interface gigabitethernet 9/3 switch(config-if)# | Configures the specified Gigabit Ethernet interface (slot 9, port 3). |
| Step 6 | switch(config-if)# channel-group 10 gigabitethernet 9/3 added to port-channel 10 and disabled please do the same operation on the switch at the other end of the port-channel, then do “no shutdown” at both ends to bring them up switch(config-if)# | Adds Gigabit Ethernet interfaces 9/3 to channel group 10. If channel group 10 does not exist, it is created. The port is shut down. |
| Step 7 | switch(config-if)# no shutdown | Enables the selected interface. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|---------|---|--|
| Step 8 | switch(config)# interface gigabitethernet 9/4 switch(config-if)# | Configures the specified Gigabit Ethernet interface (slot 9, port 4). |
| Step 9 | switch(config-if)# channel-group 10 gigabitethernet 9/4 added to port-channel 10 and disabled please do the same operation on the switch at the other end of the port-channel, then do "no shutdown" at both ends to bring them up | Adds Gigabit Ethernet interfaces 9/4 to channel group 10. The port is shut down. |
| Step 10 | switch(config-if)# no shutdown | Enables the selected interface. |



Note

Gigabit Ethernet interfaces cannot be added to a PortChannel if one of the following cases apply:

- The interface already has an IP address assigned.
- The subinterfaces are configured on that interface.
- The interface already has an associated IP-ACL rule and the PortChannel does not.

Configuring CDP


The Cisco Discovery Protocol (CDP) is supported on the management Ethernet interface on the supervisor module and the Gigabit Ethernet interfaces on the IPS module or MPS-14/2 module.

See the “Configuring CDP” section on page 4-40.

Configuring FCIP

The Fibre Channel over IP Protocol (FCIP) is a tunneling protocol that connects geographically distributed Fibre Channel storage area networks (SAN islands) transparently over IP local area networks (LANs), metropolitan area networks (MANs), and wide area networks (WANs). See Figure 28-6.

Figure 28-6 Fibre Channel SANs Connected by FCIP



FCIP uses TCP as a network layer transport.



Note

For more information about FCIP protocols, refer to the IETF standards for IP storage at <http://www.ietf.org>. Also refer to Fibre Channel standards for switch backbone connection at <http://www.t11.org> (see FC-BB-2).

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure IPS modules or MPS-14/2 modules for FCIP, you should have a basic understanding of the following concepts:

- [FCIP and VE Ports, page 28-20](#)
- [FCIP Links, page 28-21](#)
- [FCIP Profiles, page 28-21](#)
- [FCIP Interfaces, page 28-22](#)

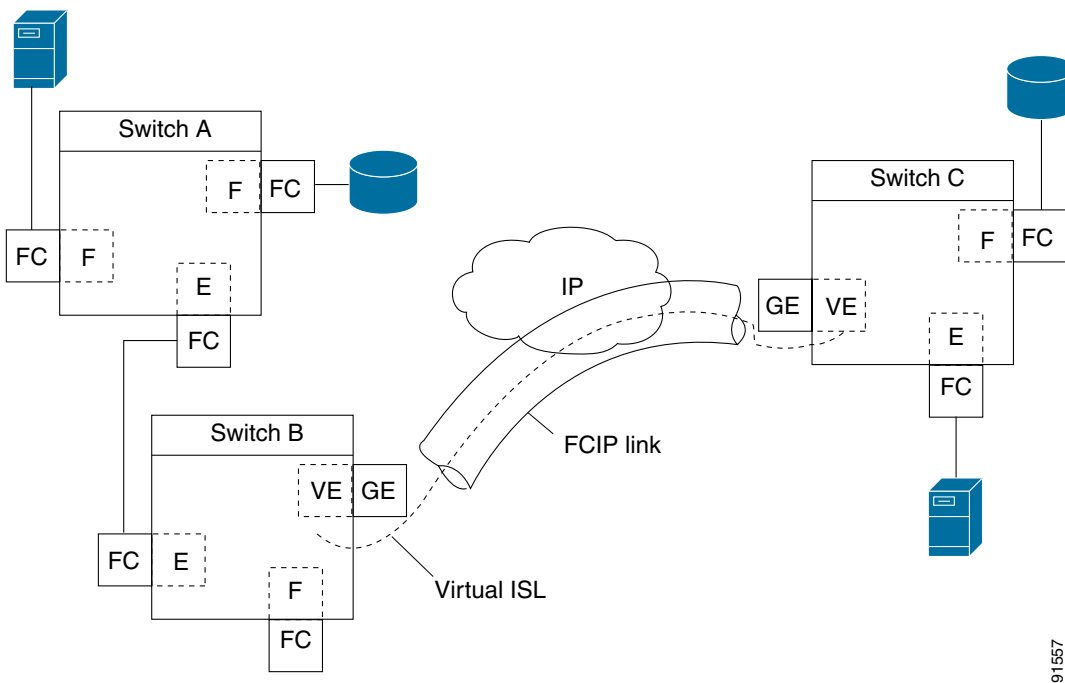
FCIP and VE Ports

[Figure 28-7](#) describes the internal model of FCIP with respect to Fibre Channel Inter-Switch Links (ISLs) and Cisco's enhanced ISLs (EISLs).

FCIP virtual E (VE) ports behave exactly like standard Fibre Channel E ports, except that the transport in this case is FCIP instead of Fibre Channel. The only requirement is for the other end of the VE port to be another VE port.

A virtual ISL is established over an FCIP link and transports Fibre Channel traffic. Each associated virtual ISL looks like a Fibre Channel ISL with either an E port or a TE port at each end (see [Figure 28-7](#)).

Figure 28-7 FCIP Links and Virtual ISLs



91557

See the “E Port” section on page 12-3.

Send documentation comments to mdsfeedback-doc@cisco.com.

FCIP Links

FCIP links consist of one or more TCP connections between two FCIP link endpoints. Each link carries encapsulated Fibre Channel frames.

When the FCIP link comes up, the VE ports at both ends of the FCIP link create a virtual Fibre Channel (E)ISL and initiate the E port protocol to bring up the (E)ISL.

By default, the FCIP feature on any Cisco MDS 9000 Family switch creates two TCP connections for each FCIP link.

- One connection is used for data frames.
- The other connection is used only for Fibre Channel control frames, that is, switch-to-switch protocol frames (all Class F). This arrangement provides low latency for all control frames.

To enable FCIP on the IPS module or MPS-14/2 module, an FCIP profile and FCIP interface (interface FCIP) must be configured.

The FCIP link is established between two peers, the VE port initialization behavior is identical to a normal E port. This behavior is independent of the link being FCIP or pure Fibre Channel, and is based on the E port discovery process (ELP, ESC).

Once the FCIP link is established, the VE port behavior is identical to E port behavior for all inter-switch communication (including domain management, zones, and VSANs). At the Fibre Channel layer, all VE and E port operations are identical.

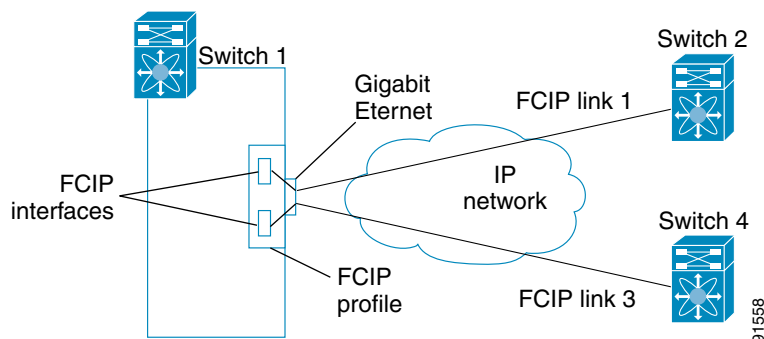
FCIP Profiles

The FCIP profile contains information about local IP address and TCP parameters. The profile defines the following information:

- The local connection points (IP address and TCP port number).
- The behavior of the underlying TCP connections for all FCIP links that use this profile.

The FCIP profile's local IP address determines the Gigabit Ethernet port where the FCIP links terminate (see [Figure 28-8](#)).

Figure 28-8 FCIP Profile and FCIP Links



Send documentation comments to mdsfeedback-doc@cisco.com.

FCIP Interfaces

The FCIP interface is the local endpoint of the FCIP link and a VE port interface. All the FCIP and E port parameters are configured in context to the FCIP interface.

The FCIP parameters consist of the following:

- The FCIP profile determines which Gigabit Ethernet port initiates the FCIP links and defines the TCP connection behavior.
- Peer information.
- Number of TCP connections for the FCIP link.
- E port parameters—trunking mode and trunk allowed VSAN list.

Enabling FCIP

To begin configuring the FCIP feature, you must explicitly enable FCIP on the required switches in the fabric. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family.

The configuration and verification commands for the FCIP feature are only available when FCIP is enabled on a switch. When you disable this feature, all related configurations are automatically discarded.

To use the FCIP feature, you need to obtain the ENTERPRISE_PKG license (see [Chapter 3, “Obtaining and Installing Licenses”](#)).

To enable FCIP on any participating switch, follow these steps:

| | Command | Purpose |
|--------|---------------------------------------|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# fcip enable | Enables FCIP on that switch. |
| | switch(config)# no fcip enable | Disables (default) FCIP on that switch. |

Basic FCIP Configuration

To configure an FCIP link, follow these steps on both switches:


| | |
|--------|---|
| Step 1 | Configure the Gigabit Ethernet interface. |
| Step 2 | Create an FCIP profile, and then assign the Gigabit Ethernet interface’s IP address to the profile. |
| Step 3 | Create an FCIP interface, and then assign the profile to the interface. |
| Step 4 | Configure the peer IP address for the FCIP interface. |
| Step 5 | Enable the interface. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Creating FCIP Profiles

You must assign a local IP address of a Gigabit Ethernet interface or subinterface to the FCIP profile to create an FCIP profile (see [Figure 28-9](#)).

Figure 28-9 Assigning Profiles to Each Gigabit Ethernet Interface



To create an FCIP profile in switch 1, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch1# config terminal switch1(config)# | Enters configuration mode. |
| Step 2 | switch1(config)# fcip profile 10 switch1(config-profile)# | Creates a profile for the FCIP connection. The valid range is from 1 to 255. |
| Step 3 | switch1(config-profile)# ip address 10.100.1.25 | Associates the profile (10) with the local IP address of the Gigabit Ethernet interface (3/1). |

To assign FCIP profile in switch 2, follow these steps:

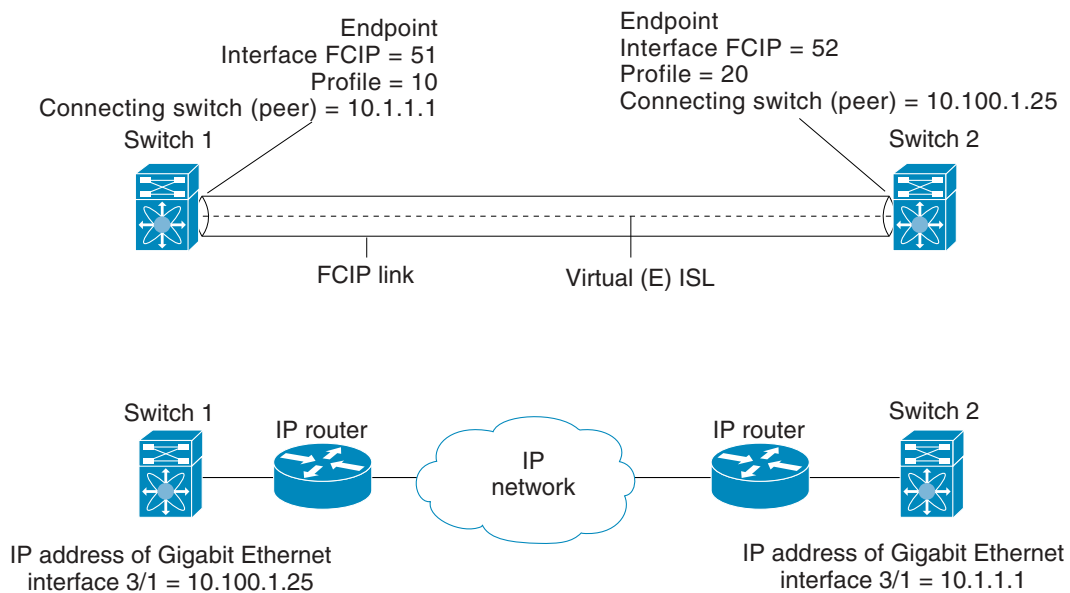
| | Command | Purpose |
|--------|---|--|
| Step 1 | switch2# config terminal switch2(config)# | Enters configuration mode. |
| Step 2 | switch2(config)# fcip profile 20 switch2(config-profile)# | Creates a profile for the FCIP connection. |
| Step 3 | switch2(config-profile)# ip address 10.1.1.1 | Associates the profile (20) with the local IP address of the Gigabit Ethernet interface. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Creating FCIP Links

When two FCIP link endpoints are created, an FCIP link is established between the two IPS modules or MPS-14/2 modules. To create an FCIP link, assign a profile to the FCIP interface and configure the peer information. The peer IP switch information initiates (creates) an FCIP link to that peer switch (see Figure 28-10).

Figure 28-10 Assigning Profiles to Each Gigabit Ethernet Interface



91562

To create an FCIP link endpoints in switch 1, follow these steps:

| | Command | Purpose |
|---------------|--|--|
| Step 1 | switch1# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch1(config)# interface fcip 51 switch1(config-if)# | Creates an FCIP interface (51). |
| Step 3 | switch1(config-if)# use-profile 10 | Assigns the profile (10) to the FCIP interface. |
| Step 4 | switch1(config-if)# peer-info ipaddr 10.1.1.1 | Assigns the peer IP address information (10.1.1.1 for switch 2) to the FCIP interface. |
| Step 5 | switch1(config-if)# no shutdown | Enables the interface. |

To create an FCIP link endpoints in switch 2, follow these steps:

| | Command | Purpose |
|---------------|--|---|
| Step 1 | switch2# config terminal switch2(config)# | Enters configuration mode. |
| Step 2 | switch2(config)# interface fcip 52 switch2(config-if)# | Creates an FCIP interface (52). |
| Step 3 | switch2(config-if)# use-profile 20 | Binds the profile (20) to the FCIP interface. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|---|---|
| Step 4 | switch2(config-if)# peer-info ip address 10.100.1.25 | Assigns the peer IP address information (10.100.1.25 for switch 1) to the FCIP interface. |
| Step 5 | switch1(config-if)# no shutdown | Enables the interface. |

Advanced FCIP Profile Configuration

A basic FCIP configuration uses the local IP address to configure the FCIP profile. In addition to the local IP address and the local port, you can specify other TCP parameters as part of the FCIP profile configuration.

- [Configuring TCP Listener Ports, page 28-25](#)
- [Configuring TCP Parameters, page 28-25](#)

FCIP configuration options can be accessed from the `switch(config-profile)#` submode prompt.

To enter the `switch(config-profile)#` prompt, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fcip profile 20 switch(config-profile)# | Creates the profile (if it does not already exist) and enters profile configuration submode. The valid range is from 1 to 255. |

Configuring TCP Listener Ports

The default TCP port for FCIP is 3225. You can change this port using the **port** command.

To change the default FCIP port number (3225), follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch(config-profile)# port 5000 | Associates the profile with the local port number (5000). |
| | switch(config-profile)# no port | Reverts to the default 3225 port. |

Configuring TCP Parameters

You can control TCP behavior in a switch by configuring the following TCP parameters.

- [Minimum Retransmit Timeout, page 28-26](#)
- [Keepalive Timeout, page 28-26](#)
- [Maximum Retransmissions, page 28-26](#)
- [Path MTUs, page 28-27](#)
- [Selective Acknowledgments, page 28-27](#)
- [Window Management, page 28-27](#)
- [Buffer Size, page 28-29](#)
- [Monitoring Congestion, page 28-28](#)
- [Estimating Maximum Jitter, page 28-29](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

Minimum Retransmit Timeout

You can control the minimum amount of time TCP waits before retransmitting. By default, this value is 200 milliseconds (ms).

To configure the minimum retransmit time, follow these steps:

| Step 1 | Command | Purpose |
|--------|---|--|
| | <code>switch(config-profile)# tcp min-retransmit-time 500</code> | Specifies the minimum TCP retransmit time for the TCP connection to be 500 ms. The default is 200 ms and the range is from 200 to 5000 ms. |
| | <code>switch(config-profile)# no tcp min-retransmit-time 500</code> | Reverts the minimum TCP retransmit time to the factory default of 200 ms. |

Keepalive Timeout

You can configure the interval that the TCP connection uses to verify that the FCIP link is functioning. This ensures that an FCIP link failure is detected quickly even when there is no traffic.

If the TCP connection is idle for more than the specified time, then keepalive timeout packets are sent to ensure that the connection is active. This command can be used to tune the time taken to detect FCIP link failures.

You can configure the first interval during which the connection is idle (the default is 60 seconds). When the connection is idle for the configured interval, eight keepalive probes are sent at 1-second intervals. If no response is received for these eight probes and the connection remains idle throughout, that FCIP link is automatically closed.



Note

Only the first interval (during which the connection is idle) can be changed.

To configure the first keepalive timeout interval, follow these steps:

| Step 1 | Command | Purpose |
|--------|---|--|
| | <code>switch(config-profile)# tcp keepalive-timeout 120</code> | Specifies the keepalive timeout interval for the TCP connection in seconds (120). The range is from 1 to 7200 seconds. |
| | <code>switch(config-profile)# no tcp keepalive-timeout 120</code> | Reverts the keepalive timeout interval to the default 60 seconds. |

Maximum Retransmissions

You can specify the maximum number of times a packet is retransmitted before TCP decides to close the connection. To configure maximum retransmissions, follow these steps:

| Step 1 | Command | Purpose |
|--------|---|--|
| | <code>switch(config-profile)# tcp max-retransmissions 6</code> | Specifies the maximum number of retransmissions (6). The range is from 1 to 8 retransmissions. |
| | <code>switch(config-profile)# no tcp max-retransmissions 6</code> | Reverts to the default of 4 retransmissions. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Path MTUs

Path MTU (PMTU) is the minimum MTU on the IP network between the two endpoints of the FCIP link. PMTU discovery is a mechanism by which TCP learns of the PMTU dynamically and adjusts the maximum TCP segment accordingly (RFC 1191).

By default, PMTU discovery is enabled on all switches with a timeout of 3600 seconds. If TCP reduces the size of the maximum segment because of PMTU change, the reset-timeout specifies the time after which TCP tries the original MTU.

To configure PMTU, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | <code>switch(config-profile)# no tcp pmtu-enable</code> | Disables PMTU discovery. |
| | <code>switch(config-profile)# tcp pmtu-enable</code> | Enables (default) PMTU discovery with the default value of 3600 seconds. |
| | <code>switch(config-profile)# tcp pmtu-enable reset-timeout 90</code> | Specifies the PMTU reset timeout to 90 seconds. The range is 60 to 3600 seconds. |
| | <code>switch(config-profile)# no tcp pmtu-enable reset-timeout 600</code> | Leaves PMTU discovery enabled but reverts the timeout to the default of 3600 seconds. |

Selective Acknowledgments

TCP may experience poor performance when multiple packets are lost within one window. With the limited information available from cumulative acknowledgments, a TCP sender can only learn about a single lost packet per round trip. A selective acknowledgment (SACK) mechanism helps overcome the limitations of multiple lost packets during a TCP transmission.

The receiving TCP sends back SACK advertisements to the sender. The sender can then retransmit only the missing data segments. By default, SACK is enabled on Cisco MDS 9000 Family switches.

To configure SACK, follow these steps:

| | Command | Purpose |
|--------|---|-------------------------|
| Step 1 | <code>switch(config-profile)# no tcp sack-enable</code> | Disables SACK. |
| | <code>switch(config-profile)# tcp sack-enable</code> | Enables SACK (default). |

Window Management

The optimal TCP window size is automatically calculated using the maximum bandwidth parameter, the minimum available bandwidth parameter, and the dynamically measured round trip time (RTT).



Note

The configured **round-trip-time** parameter determines the window scaling factor of the TCP connection. This parameter is only an approximation. The measured RTT value overrides the round trip time parameter for window management. If the configured **round-trip-time** is too small compared to the measured RTT, then the link may not be fully utilized due to the window scaling factor being too small.

The **min-available-bandwidth** parameter and the measured RTT together determine the threshold below which TCP aggressively maintains a window size sufficient to transmit at minimum available bandwidth.

The **max-bandwidth-mbps** parameter and the measured RTT together determine the maximum window size.

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure window management, follow these steps:

| Step 1 | Command | Purpose |
|--------|---|---|
| | <code>switch(config-profile)# tcp max-bandwidth-mbps 900 min-available-bandwidth-mbps 300 round-trip-time-ms 10</code> | Configures the maximum available bandwidth at 900 Mbps, the minimum slow start threshold at 300 Mbps, and the RTT at 10 ms. |
| | <code>switch(config-profile)# no tcp max-bandwidth-mbps 900 min-available-bandwidth-mbps 300 round-trip-time-ms 10</code> | Reverts to the factory defaults. The FCIP defaults are maximum bandwidth at 1 Gbps, minimum available bandwidth at 500 Mbps, and RTT at 1 ms. |
| | <code>switch(config-profile)# tcp max-bandwidth-kbps 2000 min-available-bandwidth-kbps 2000 round-trip-time-us 200</code> | Configures the maximum available bandwidth at 2000 Kbps, the minimum available bandwidth at 2000 Kbps, and the RTT at 200 ms. |

Monitoring Congestion

By enabling the congestion window monitoring (CWM) parameter, you allow TCP to monitor congestion after each idle period. The **cwm** parameter also determines the maximum burst size allowed after an idle period. By default, this parameter is enabled and the default burst size is 50 KB.

The interaction of bandwidth parameters and CWM and the resulting TCP behavior is outlined below:

- If the average rate of the Fibre Channel traffic over the preceding RTT is less than the min-available-bandwidth multiplied by the RTT, the entire burst is sent immediately at the min-available-bandwidth rate, provided no TCP drops occur.
- If the average rate of the Fibre Channel traffic is greater than min-available-bandwidth multiplied by the RTT, but less than max-bandwidth multiplied by the RTT, then if the Fibre Channel traffic is transmitted in burst sizes smaller than the configured CWM value the entire burst is sent immediately by FCIP at the max-bandwidth rate.
- If the average rate of the Fibre Channel traffic is larger than the min-available-bandwidth multiplied by the RTT and the burst size is greater than the CWM value, then only a part of the burst is sent immediately. The remainder is sent with the next RTT.

The software uses standard TCP rules to increase the window beyond the one required to maintain the min-available-bandwidth to reach the max-bandwidth.



Note

As of Cisco SAN-OS Release 2.0(1b), the default burst size is 50 KB. In Cisco SAN-OS Release 1.3(1) and earlier, the burst size was 10KB.



Tip

We recommend that this feature remain enabled to realize optimal performance. Increasing the CWM burst size can result in more packet drops in the IP network, impacting TCP performance. Only if the IP network has sufficient buffering, try increasing the CWM burst size beyond the default to achieve lower transmit latency.

Send documentation comments to mdsfeedback-doc@cisco.com.

To change the CWM defaults, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | <code>switch(config-profile)# no tcp cwm</code> | Disables congestion monitoring. |
| | <code>switch(config-profile)# tcp cwm</code> | Enables congestion monitoring and sets the burst size to its default size. |
| | <code>switch(config-profile)# tcp cwm burstsize 30</code> | Changes the burst size to 30 KB. The valid range is from 10 to 100 KB. |
| | <code>switch(config-profile)# no tcp cwm burstsize 25</code> | Leaves the CWM feature in an enabled state but changes the burst size to its factory default. |

Estimating Maximum Jitter

Jitter is defined as a variation in the delay of received packets. At the sending side, packets are sent in a continuous stream with the packets spaced evenly apart. Due to network congestion, improper queuing, or configuration errors, this steady stream can become lumpy, or the delay between each packet can vary instead of remaining constant.

As of Cisco SAN-OS Release 1.3(4a), you can configure the maximum estimated jitter in microseconds by the packet sender. The estimated variation should not include network queuing delay. By default, this parameter is enabled in Cisco MDS switches when IPS modules or MPS-14/2 modules are present.

The default value is 1000 microseconds for FCIP interfaces and 500 microseconds for iSCSI interfaces.



Note

As of Cisco SAN-OS Release 2.0(1b), the default value for FCIP interfaces is 1000 microseconds. In Cisco SAN-OS Release 1.3(4a) (and later), the burst size was 100 microseconds.

To configure the maximum jitter value, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | <code>switch(config-profile)# no tcp max-jitter</code> | Disables delay jitter estimation. |
| | <code>switch(config-profile)# tcp max-jitter</code> | Enables the delay jitter feature and sets the time to its factory default. |
| | <code>switch(config-profile)# tcp max-jitter 300</code> | Changes the time to 300 microseconds. The valid range is from 0 to 10000 microseconds. |
| | <code>switch(config-profile)# no tcp max-jitter 2500</code> | Leaves the delay jitter feature in an enabled state but changes the time to its factory default. |

Buffer Size

You can define the required additional buffering—beyond the normal send window size—that TCP allows before flow controlling the switch's egress path for the FCIP interface. The default FCIP buffer size is 0 KB.



Note

Use the default if the FCIP traffic is passing through a high throughput WAN link. If you have a mismatch in speed between the Fibre Channel link and the WAN link, then time stamp errors occur in the DMA bridge. In such a situation, you can avoid time stamp errors by increasing the buffer size.

Send documentation comments to mdsfeedback-doc@cisco.com.

To set the buffer size, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | <code>switch(config-profile)# tcp send-buffer-size 5000</code> | Configure the advertised buffer size to 5000 KB. As of Cisco SAN-OS Release 2.0(1b), the valid range is from 0 to 16384 KB. |
| | <code>switch(config-profile)# no tcp send-buffer-size 5000</code> | Reverts the switch to its factory default. The default is 0 KB. |

Advanced FCIP Interface Configuration

You can establish connection to a peer by configuring one or more of the following options for the FCIP interface.

- [Configuring Peers, page 28-30](#)
- [Active Connections, page 28-32](#)
- [Number of TCP Connections, page 28-32](#)
- [Time Stamp Control, page 28-32](#)
- [B Port Interoperability Mode, page 28-34](#)
- [Quality of Service, page 28-36](#)

To establish a peer connection, you must first create the FCIP interface and enter the `config-if` submode.

To enter the `config-if` submode, follow these steps:

| | Command | Purpose |
|--------|---|----------------------------------|
| Step 1 | <code>switch# config terminal</code> | Enters configuration mode. |
| Step 2 | <code>switch(config)# interface fcip 100</code> | Creates an FCIP interface (100). |

Configuring Peers

To establish an FCIP link with the peer, you can use one of two options:

- Peer IP address—Configures both ends of the FCIP link. Optionally, you can also use the peer TCP port along with the IP address.
- Special frames—Configures one end of the FCIP link when security gateways are present in the IP network. Optionally, you can also use the switch WWN (sWWN) and profile ID along with the IP address.

Send documentation comments to mdsfeedback-doc@cisco.com.

Peer IP Address

The basic FCIP configuration uses the peer's IP address to configure the peer information. You can also specify the peer's port number to configure the peer information. If you do not specify a port, the default 3225 port number is used to establish connection.

To assign the peer information based on the IP address, port number, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch(config-if)# peer-info ipaddr 10.1.1.1 | Assigns an IP address to configure the peer information. Because no port is specified, the default port number, 3225, is used. |
| | switch(config-if)# no peer-info ipaddr 10.10.1.1 | Deletes the assigned peer port information. |
| Step 2 | switch(config-if)# peer-info ipaddr 10.1.1.1 port 3000 | Assigns the IP address and sets the peer TCP port to 3000. The valid port number range is from 0 to 65535. |
| | switch(config-if)# no peer-info ipaddr 10.1.1.1 port 2000 | Deletes the assigned peer port information. |
| Step 3 | switch(config-if)# no shutdown | Enables the interface. |

Special Frames

You can alternatively establish an FCIP link with a peer using an optional protocol called special frames. When special frames are enabled, the peer IP address (and optionally the port or the profile ID) only needs to be configured on one end of the link. Once the connection is established, a special frame is exchanged to discover and authenticate the link.

By default, the special frame feature is disabled.



Note

Refer to the Fibre Channel IP standards for further information on special frames.



Tip

Special frame negotiation provides an additional authentication security mechanism because the link validates the WWN of the peer switch.

You can enable or disable the **special-frame** option. On both sides, the **special-frame** option must be enabled to establish the FCIP link.

To enable special frames, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch(config-if)# special-frame peer-wnn 12:12:34:45:ab:bc:cd:00 | Enables special frames and sets the peer WWN as specified. Note The peer WWN is the WWN of the peer switch. Use the show wwn switch command to obtain the peer WWN. |
| | switch(config-if)# no special-frame peer-wnn 12:12:34:45:ab:bc:cd:00 | Disables special frames (default). |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|--|--|
| Step 2 | <code>switch(config-if)# special-frame peer-wwn 12:12:34:45:ab:bc:cd:00 peer profile-id 155</code> | Enables special frames and sets the peer WWN and the profile ID (155). |
| | <code>switch(config-if)# no special-frame peer-wwn 12:12:34:45:ab:bc:cd:00 peer profile-id 155</code> | Disables special frames (default). |
| Step 3 | <code>switch(config-if)# no shutdown</code> | Enables the interface. |

Active Connections

You can configure the required mode for initiating an TCP connection. By default, active mode is enabled to actively attempt an IP connection. If you enable the passive mode, the switch does not initiate a TCP connection and merely waits for the peer to connect to it.



Note Ensure that both ends of the FCIP link are not configured as passive mode. If both ends are configured as passive, the connection is not initiated.

To enable the passive mode, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | <code>switch(config-if)# passive-mode</code> | Enables passive mode while attempting a TCP connection. |
| | <code>switch(config-if)# no passive-mode</code> | Reverts to the factory set default of using the active mode while attempting the TCP connection. |
| Step 2 | <code>switch(config-if)# no shutdown</code> | Enables the interface. |

Number of TCP Connections

You can specify the number of TCP connections from an FCIP link. By default, the switch tries two (2) TCP connections for each FCIP link. You can configure one or two TCP connections. For example, the Cisco PA-FC-1G Fibre Channel port adapter, which has only one (1) TCP connection, interoperates with any switch in the Cisco MDS 9000 Family. One TCP connection is within the specified limit. If the peer initiates one TCP connection, and your MDS switch is configured for two TCP connections, then the software handles it gracefully and moves on with just one connection.

To specify the TCP connection attempts, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | <code>switch(config-if)# tcp-connection 1</code> | Specifies the number of TCP connections. Valid values are 1 or 2. |
| | <code>switch(config-if)# no tcp-connection 1</code> | Reverts to the factory set default of two attempts. |
| Step 2 | <code>switch(config-if)# no shutdown</code> | Enables the interface. |

Time Stamp Control

You can instruct the switch to discard packets that are outside the specified time. By default, this option is disabled in all switches in the Cisco MDS 9000 Family. The **acceptable-diff** option specifies the time range within which packets can be accepted. If the packet arrived within the range specified by this option, the packet is accepted. Otherwise, it is dropped. By default, if a packet arrives within a 2000 millisecond interval (+ or -2000 ms) from the network time, that packet is accepted.

Send documentation comments to mdsfeedback-doc@cisco.com.

**Note**

As of Cisco SAN-OS Release 2.0(1b), the default value for packet acceptance is 2000 microseconds. In Cisco SAN-OS Release 1.3(1) and earlier, the burst size was 1000 microseconds.

**Note**

If the **time-stamp** option is enabled, be sure to configure NTP on both switches (see the “[NTP Configuration](#)” section on page 4-17).

**Tip**

Do not enable time stamp control on an FCIP interface that has tape acceleration or write acceleration configured.

To enable or disable the time stamp control, follow these steps:

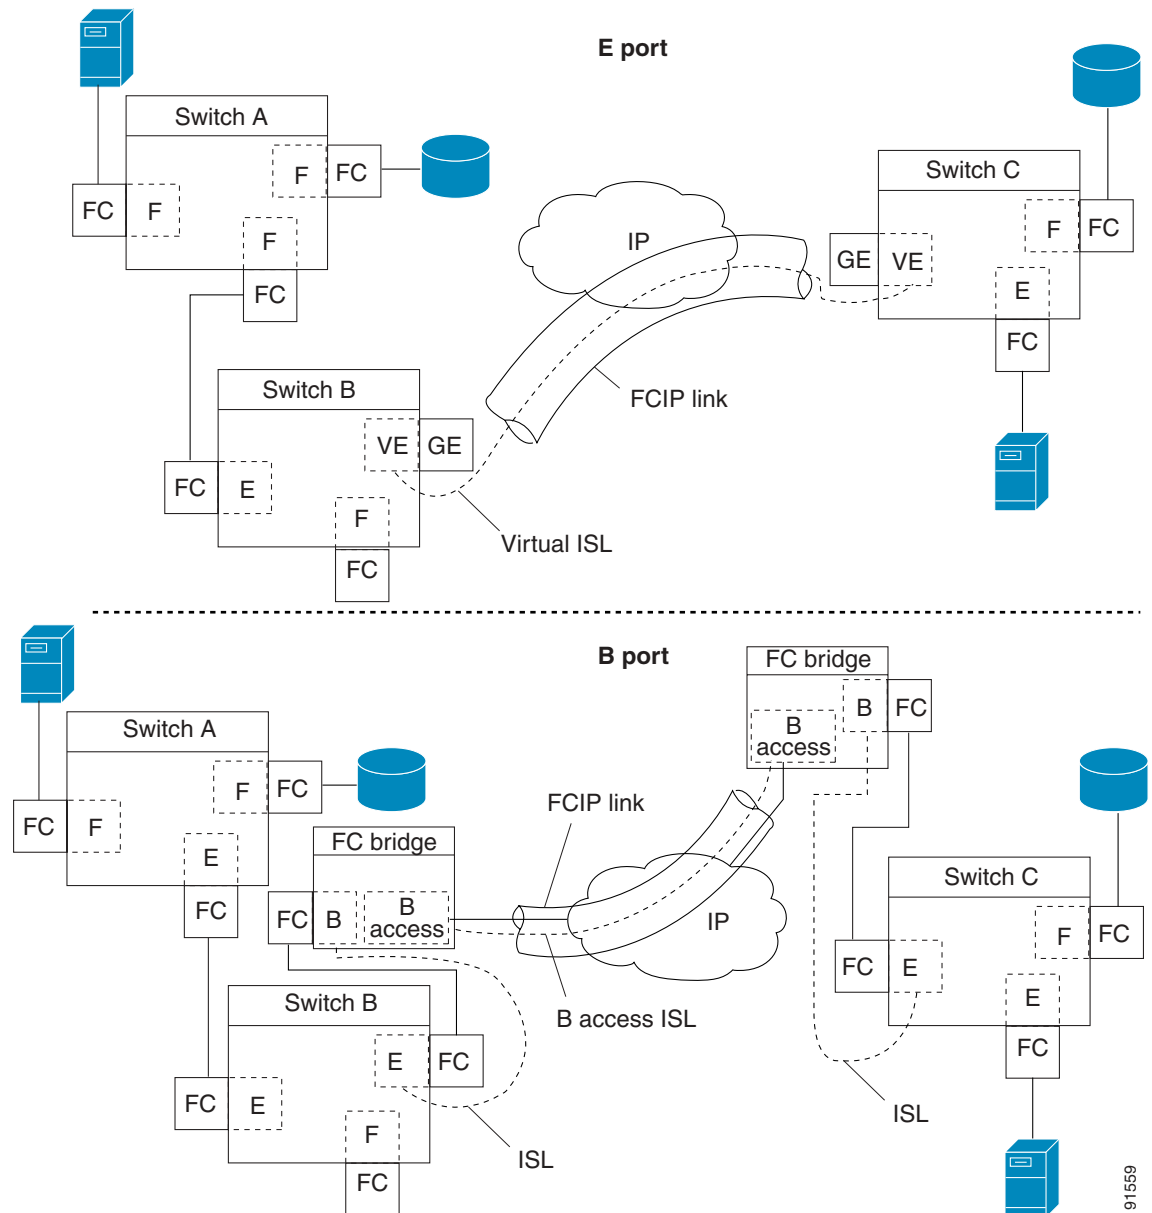
| | Command | Purpose |
|---------------|---|---|
| Step 1 | switch(config-if)# time-stamp Please enable NTP with a common time source on both MDS Switches that are on either side of the FCIP link | Enables time stamp checking for received packets with a default acceptable time difference of 2000 ms. |
| | switch(config-if)# no time-stamp | Disables (default) time stamps. |
| Step 2 | switch(config-if)# time-stamp acceptable-diff 4000 | Configures the packet acceptance time. The valid range is from 500 to 10,000 ms. |
| | switch(config-if)# no time-stamp acceptable-diff 500 | Deletes the configured time difference and reverts the difference to factory defaults. The default difference is a 2000 millisecond interval from the network time. |
| Step 3 | switch(config-if)# no shutdown | Enables the interface. |

Send documentation comments to mdsfeedback-doc@cisco.com.

B Port Interoperability Mode

While E ports typically interconnect Fibre Channel switches, some SAN extender devices, such as Cisco's PA-FC-1G Fibre Channel port adapter and the SN 5428-2 storage router, implement a bridge port model to connect geographically dispersed fabrics. This model uses B port as described in the T11 Standard FC-BB-2. [Figure 28-11](#) depicts a typical SAN extension over an IP network.

Figure 28-11 FCIP B Port and Fibre Channel E Port



B ports bridge Fibre Channel traffic from a local E port to a remote E port without participating in fabric-related activities such as principal switch election, domain ID assignment, and Fibre Channel fabric shortest path first (FSPF) routing. For example, Class F traffic entering a SAN extender does not

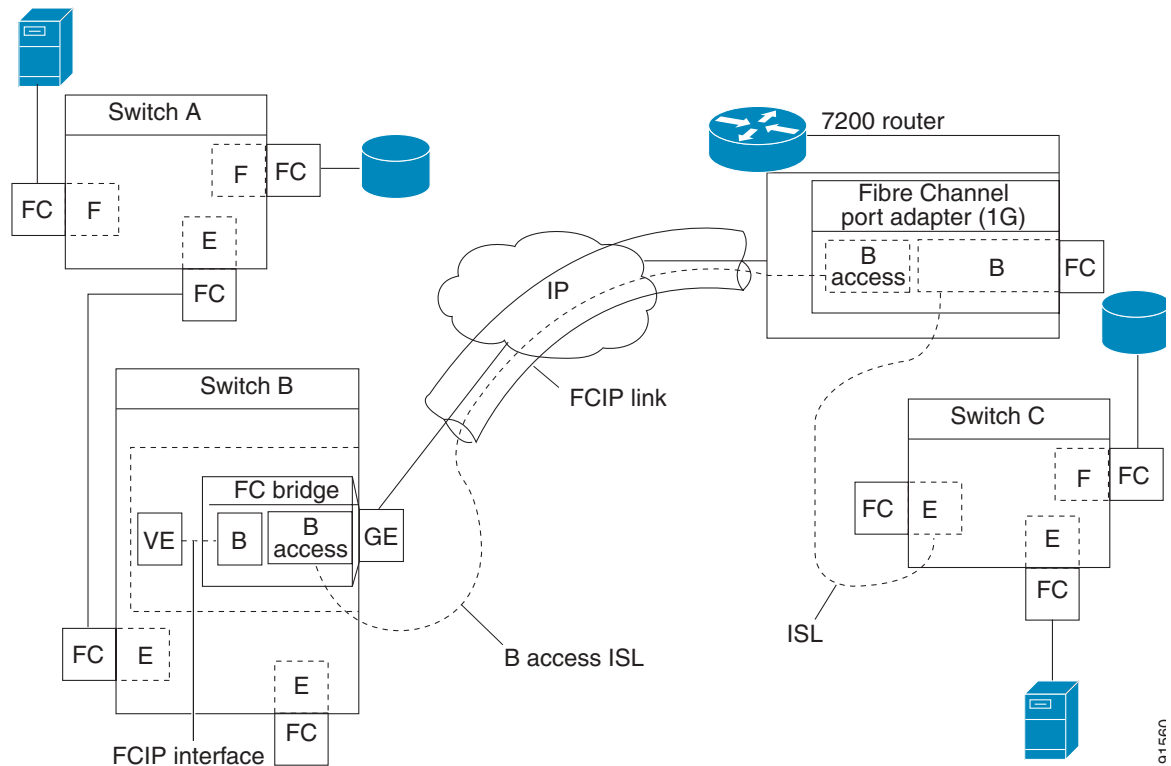
Send documentation comments to mdsfeedback-doc@cisco.com.

interact with the B port. The traffic is transparently propagated (bridged) over a WAN interface before exiting the remote B port. This bridge results in both E ports exchanging Class F information that ultimately leads to normal ISL behavior such as fabric merging and routing.

FCIP links between B port SAN extenders do not exchange the same information as FCIP links between E ports, and are therefore incompatible. This is reflected by the terminology used in FC-BB-2: *while VE ports establish a virtual ISL over an FCIP link, B ports use a B access ISL*.

The IPS module and MPS-14/2 module support FCIP links that originate from a B port SAN extender device by implementing the B access ISL protocol on a Gigabit Ethernet interface. Internally, the corresponding virtual B port connects to a virtual E port that completes the end-to-end E port connectivity requirement (see [Figure 28-12](#)).

Figure 28-12 FCIP Link Terminating in a B Port Mode



The B port feature in the IPS module and MPS-14/2 module allows remote B port SAN extenders to communicate directly with a Cisco MDS 9000 Family switch, therefore eliminating the need for local bridge devices.

Send documentation comments to mdsfeedback-doc@cisco.com.

Configuring B Ports

When an FCIP peer is a SAN extender device that only supports Fibre Channel B ports, you need to enable the B port mode for the FCIP link. When a B port is enabled, the E port functionality is also enabled and they coexist. If the B port is disabled, the E port functionality remains enabled.

To enable B port mode, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | <code>switch(config-if)# bport</code> | Enables B port mode on the FCIP interface. |
| | <code>switch(config-if)# no bport</code> | Reverts to E port mode on the FCIP interface (default). |
| Step 2 | <code>switch(config-if)# bport-keepalive</code> | Enables the reception of keepalive responses sent by a remote peer. |
| | <code>switch(config-if)# no bport-keepalive</code> | Disables the reception of keepalive responses sent by a remote peer (default). |

Quality of Service

The Quality of Service (QoS) parameter specifies the differentiated services code point (DSCP) value to mark all IP packets (type of service—TOS field in the IP header).

- The control DSCP value applies to all FCIP frames in the control TCP connection.
- The data DSCP value applies to all FCIP frames in the data connection.

If the FCIP link has only one TCP connection, that data DSCP value is applied to all packets in that connection.

To set the QoS values, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | <code>switch(config-if)# qos control 3 data 5</code> | Configures the control TCP connection and data connection to mark all packets on that DSCP value. The control and data value ranges from 0 to 63. |
| | <code>switch(config-if)# no qos control 3 data 5</code> | Reverts the switch to its factory default (marks all packets with DSCP value 0). |

Configuring E Ports

All configuration commands that apply to E ports, also apply to FCIP interfaces. The following features are also available for FCIP interfaces:

- An FCIP interface can be a member of any VSAN (see [Chapter 10, “Configuring and Managing VSANs”](#)).
- Trunk mode and trunk allowed VSANs (see [Chapter 13, “Configuring Trunking”](#)).
- PortChannels (see [Chapter 14, “Configuring PortChannels”](#)):
 - Multiple FCIP links can be bundled into a Fibre Channel PortChannel.
 - FCIP links and Fibre Channel links cannot be combined in one PortChannel.
- FSPF (see [Chapter 24, “Configuring Fibre Channel Routing Services and Protocols”](#)).

Send documentation comments to mdsfeedback-doc@cisco.com.

- Fibre Channel domains (fcdomains) (see [Chapter 31, “Configuring Domain Parameters.”](#)).
- Importing and exporting the zone database from the adjacent switch (see [Chapter 15, “Configuring and Managing Zones”](#)).

Advanced FCIP Features

You can significantly improve application performance by configuring one or more of the following options for the FCIP interface.

- [FCIP Write Acceleration, page 28-37](#)
- [FCIP Tape Acceleration, page 28-39](#)
- [FCIP Compression, page 28-41](#)

FCIP Write Acceleration

The FCIP write acceleration feature in Cisco SAN-OS Release 1.3(3) (and later) enables you to significantly improve application write performance when storage traffic is routed over wide area networks using FCIP. When FCIP write acceleration is enabled, WAN throughput is maximized by minimizing the impact of WAN latency for write operations.



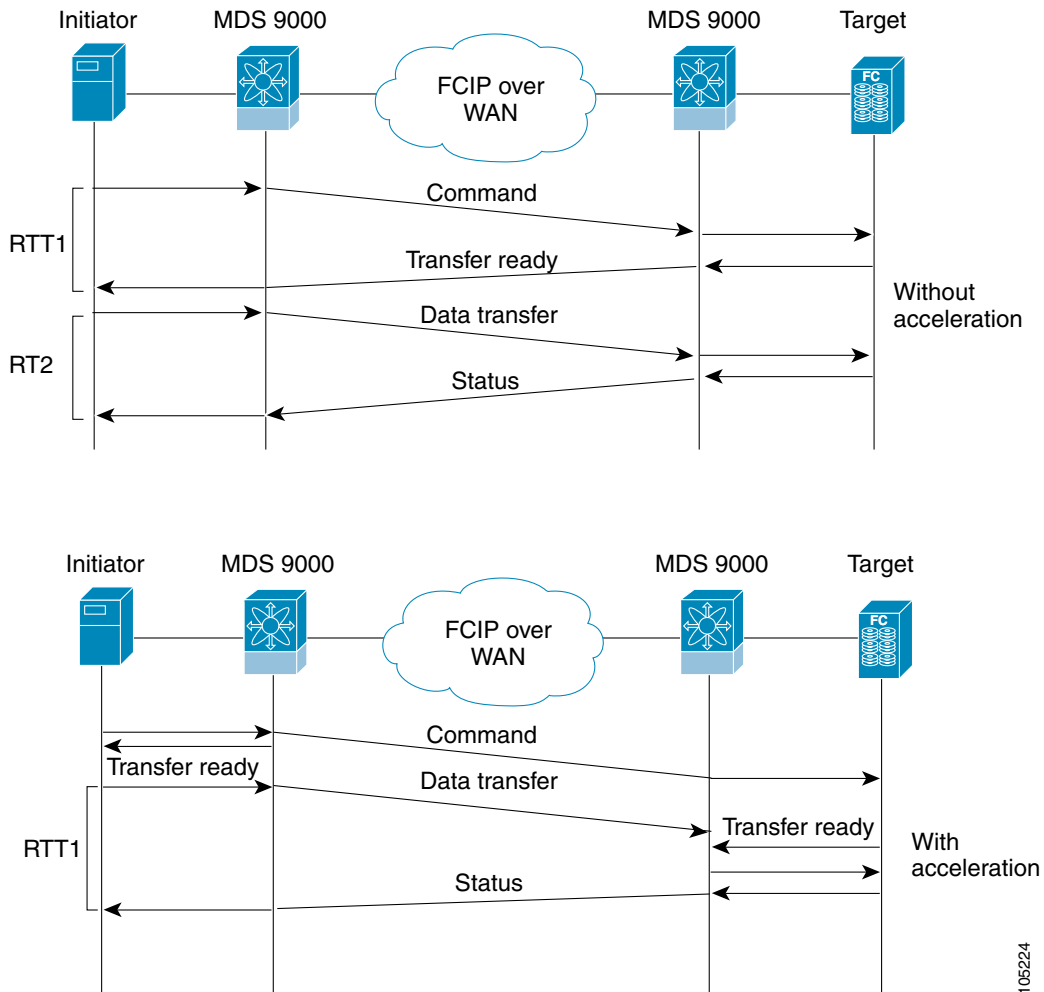
Note

The write acceleration feature is disabled by default and must be enabled on both sides of the FCIP link. If it is only enabled on one side of the FCIP tunnel, the tunnel is not initialized.

In [Figure 28-13](#), the WRITE command without write acceleration requires two round trip transfers (RTT), while the WRITE command with write acceleration only requires one RTT. The maximum sized Transfer Ready is sent from the host side of the FCIP link back to the host before the WRITE command reaches the target. This enables the host to start sending the write data without waiting for the long latency over the FCIP link of the WRITE command and Transfer Ready. It also eliminates the delay caused by multiple Transfer Readys needed for the exchange going over the FCIP link.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 28-13 FCIP Link Write Acceleration



Tip

As of Cisco SAN-OS Release 2.0(1b), FCIP write acceleration works even if the FCIP port is part of a PortChannel. In releases prior to Cisco SAN-OS Release 2.0(1b), FCIP write acceleration does not work if the FCIP port is part of a PortChannel. FCIP write acceleration also does not work if multiple paths exist with equal weight between the initiator and the target port. Such a configuration might cause either SCSI discovery failure or broken WRITE or READ operations.



Tip

Do not enable time stamp control on an FCIP interface write acceleration configured.



Caution

When write acceleration is enabled in an FCIP interface, a FICON VSAN cannot be enabled in that interface. Likewise, if an FCIP interface is up in a FICON VSAN, write acceleration cannot be enabled on that interface.

105224

Send documentation comments to mdsfeedback-doc@cisco.com.



Caution

FCIP write acceleration with FCIP ports as members of PortChannels in Cisco SAN-OS Release 2.0(1b) (or later) are incompatible with the FCIP write acceleration in earlier releases.

To enable write acceleration, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch1# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch1(config)# interface fcip 51 switch1(config-if)# | Creates an FCIP interface (51). |
| Step 3 | switch1(config-if)# write-accelerator switch1(config-if)# no write-accelerator | Enables write acceleration. Disables write acceleration (default). |

FCIP Tape Acceleration

Tapes are storage devices that store and retrieve user data sequentially. Applications that access tape drives normally have only one SCSI WRITE operation outstanding to it. This single command process limits the benefit of the write acceleration feature when using an FCIP tunnel over a long-distance WAN link. It impacts backup and archive performance because each SCSI WRITE operation does not complete until the host receives a good status response from the tape drive.

The FCIP tape acceleration feature is introduced in Cisco SAN-OS Release 2.0(1b) to help solve this problem. It improves tape backup and archive operations by allowing faster data streaming from the host to the tape over the WAN link.

The backup server in [Figure 28-14](#) issues write operations to a drive in the tape library. Acting as a proxy for the remote tape drives, the local Cisco MDS switch proxies a transfer ready to signal the host to start sending data. After receiving all the data, the local Cisco MDS switch proxies the successful completion of the SCSI WRITE operation. This response allows the host to start the next SCSI WRITE operation. This proxy method results in more data being sent over the FCIP tunnel in the same time period compared to the time taken to send data without proxying. The proxy method improves the use of WAN links.

At the tape end of the FCIP tunnel, another Cisco MDS switch buffers the command and data it has received. It then acts as a backup server to the tape drive by listening to a transfer ready from the tape drive before forwarding the data.

The Cisco SAN-OS provides reliable data delivery to the remote tape drives using TCP/IP over the WAN. It maintains write data integrity by allowing the WRITE FILEMARKS operation to complete end-to-end without proxying. The WRITE FILEMARKS operation signals the synchronization of the buffer data with the tape library data. While tape media errors are returned to backup servers for error handling, tape busy errors are retried automatically by the Cisco SAN-OS software.

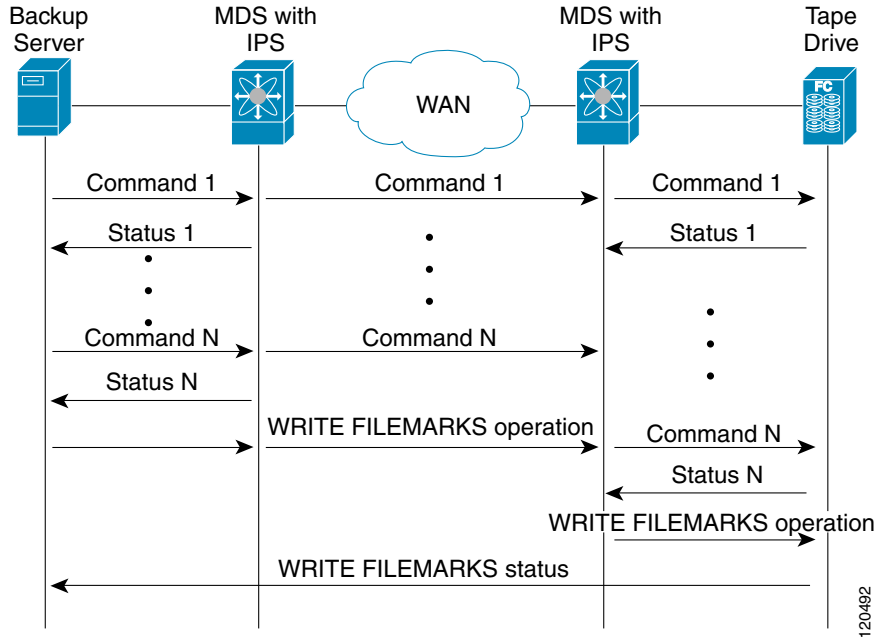


Note

The tape acceleration feature is disabled by default and must be enabled on both sides of the FCIP link. If it is only enabled on one side of the FCIP tunnel, the tunnel is not initialized.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 28-14 FCIP Link Tape Acceleration



Tip

FCIP tape acceleration does not work if the FCIP port is part of a PortChannel or if there are multiple paths with equal weight between the initiator and the target port. Such a configuration might cause either SCSI discovery failure or broken write or read operations.



Caution

When tape acceleration is enabled in an FCIP interface, a FICON VSAN cannot be enabled in that interface. Likewise, if an FCIP interface is up in a FICON VSAN, write acceleration cannot be enabled on that interface.

When you enable the tape acceleration feature for an FCIP tunnel, the tunnel is reinitialized and the write acceleration feature is also automatically enabled.

In Tape Acceleration after a certain amount of data has been buffered at the remote Cisco MDS switch, the write operations from the host are flow controlled by the local Cisco MDS switch, by not proxying the Transfer Ready. On completion of a write operation when some data buffers are freed, the local Cisco MDS switch resumes the proxying.

The default flow control buffering uses the **automatic** option. This option takes the WAN latencies and the speed of the tape into account to provide optimum performance. You can also specify a flow control buffer size (the maximum buffer size is 12MB).



Tip

We recommend that you use the default option for flow control buffering.



Tip

Do not enable time stamp control on an FCIP interface tape acceleration configured.

Send documentation comments to mdsfeedback-doc@cisco.com.

To enable tape acceleration, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch1# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch1(config)# interface fcip 5 switch1(config-if)# | Creates an FCIP interface (5). |
| Step 3 | switch1(config-if)# write-accelerator tape-accelerator | Enables tape acceleration (and write acceleration—if not already enabled). |
| | switch1(config-if)# write-accelerator tape-accelerator flow-control-buffer-size auto | Enables tape acceleration with automatic flow control (default) |
| | switch1(config-if)# write-accelerator tape-accelerator flow-control-buffer-size 2048 | Sets tape acceleration flow control buffer size to 2 MB. |
| | switch1(config-if)# no write-accelerator tape-accelerator | Disables tape acceleration (default) and resets the FCIP tunnel. Note The write-acceleration feature remains enabled. |
| | switch1(config-if)# no write-accelerator tape-accelerator flow-control-buffer-size 2048 | Changes the flow control buffer size to the default value of automatic. The tape acceleration and write acceleration features remain enabled. This command does not reset the FCIP tunnel. |
| | switch1(config-if)# no write-accelerator | Disables both the write acceleration and tape-acceleration features and resets the FCIP tunnel. |

FCIP Compression

The FCIP compression feature introduced in Cisco SAN-OS Release 1.3(1) allows IP packets to be compressed on the FCIP link if this feature is enabled on that link. By default the FCIP compression is disabled. When enabled, the software defaults to using the **auto** mode (if a mode is not specified).

As of Cisco SAN-OS Release 2.0(1b), you can configure FCIP compression using one of the following modes:

- **mode1** is a fast compression mode for high bandwidth links (> 25 Mbps)
- **mode2** is a moderate compression mode for moderately low bandwidth links (between 10 and 25 Mbps)
- **mode3** is a high compression mode for low bandwidth links (< 10 Mbps)
- **auto** (default) mode picks the appropriate compression scheme based on the bandwidth of the link (the bandwidth of the link configured in the FCIP profile's TCP parameters)

The IP compression feature behavior differs between the IPS module(s) or MPS-14/2 module(s) and MPS-14/2 module—while **mode2** and **mode3** perform software compression in both modules, **mode1** performs hardware-based compression in MPS-14/2 modules, and software compression in IPS-4 and IPS-8 modules.



Note

The Cisco MDS 9216i Switch also supports IP compression feature. The integrated supervisor module has the same hardware components that are available in the MPS-14/2 module.

Send documentation comments to mdsfeedback-doc@cisco.com.



Caution

The compression modes in Cisco SAN-OS Release 2.0(1b) and later are compatible with the compression modes in Cisco SAN-OS Release 1.3(1) and earlier.

If one end of the FCIP link is running Cisco SAN-OS Release 2.0(1b) (or later) and the other end is running Cisco SAN-OS Release 1.3(1) (or earlier), then you must disable compression at both ends of the FCIP link.



Tip

While upgrading from Cisco SAN-OS Release 1.x to Cisco SAN-OS Release 2.0(1b) or later, we recommend that you disable compression before the upgrade procedure, and then enable the required mode after the upgrade procedure.

If both ends of the FCIP link are running Cisco SAN-OS Release 2.0(1b) (or later) and you enable compression at one end of the FCIP tunnel, then be sure to enable it at the other end of the link.

To enable FCIP compression, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# interface fcip 51 switch(config-if)# | Creates an FCIP interface (51). |
| Step 3 | switch(config-if)# ip-compression mode3 | Enables high compression for low bandwidth links. |
| | switch(config-if)# ip-compression mode3 | Defaults to using the auto mode. |
| | switch(config-if)# no ip-compression | Disables (default) the FCIP compression feature. |

Displaying FCIP Information

Use the **show interface** commands to view the summary, counter, description, and status of the FCIP link. Use the output of these commands to verify the administration mode, the interface status, the operational mode, the related VSAN ID, and the profile used. See [Example 28-19](#) through [Example 28-28](#).

Example 28-14 Displays the FCIP Summary

```
switch# show fcip summary
```

| Tun | prof | Eth-if | peer-ip | Status | T | W | T | Enc | Comp | Bandwidth | rtt |
|-----|------|-----------|----------|--------|---|---|---|-----|------|-------------|------|
| | | | | | E | A | A | | | max/min | (us) |
| 10 | 91 | GE4/1 | 3.3.3.2 | UP | N | N | N | N | N | 1000M/1000M | 2000 |
| 11 | 11 | GE3/1.601 | 30.1.1.2 | DOWN | N | N | N | N | N | 1000M/500M | 1000 |
| 12 | 12 | GE3/1.602 | 30.1.2.2 | DOWN | N | N | N | N | N | 1000M/500M | 1000 |
| 13 | 0 | | 0.0.0.0 | DOWN | N | N | N | N | N | | |
| 14 | 0 | | 0.0.0.0 | DOWN | N | N | N | N | N | | |
| 15 | 0 | | 0.0.0.0 | DOWN | N | N | N | N | N | | |
| 16 | 0 | | 0.0.0.0 | DOWN | N | N | N | N | N | | |
| 17 | 0 | | 0.0.0.0 | DOWN | N | N | N | N | N | | |
| 18 | 0 | | 0.0.0.0 | DOWN | N | N | N | N | N | | |

Send documentation comments to mdsfeedback-doc@cisco.com.

```

19 0          0.0.0.0          DOWN N N N N N
20 92 GE4/2    3.3.3.1          UP   N N N N N 1000M/1000M 2000
21 21 GE3/2.601 30.1.1.1       DOWN N N N N N 1000M/500M 1000
22 22 GE3/2.602 30.1.2.1       DOWN N N N N N 1000M/500M 1000

```

Example 28-15 Displays Exchanges Processed by Write Acceleration or Tape Acceleration at the Specified Host End FCIP Link.

```

switch# show fcip host-map 100

MAP TABLE (5 entries TOTAL entries 5)

OXID | RXID | HOST FCID | TARG FCID | VSAN | Index
-----+-----+-----+-----+-----+-----
0xd490|0xffff|0x00690400|0x00620426|0x0005|0x0000321f
0xd4a8|0xffff|0x00690400|0x00620426|0x0005|0x00003220
0xd4c0|0xffff|0x00690400|0x00620426|0x0005|0x00003221
0xd4d8|0xffff|0x00690400|0x00620426|0x0005|0x00003222
0xd4f0|0xffff|0x00690400|0x00620426|0x0005|0x00003223

```

Example 28-16 Displays Exchanges Processed by Write Acceleration or Tape Acceleration at the Specified Target End FCIP Link

```

switch# show fcip target-map 100

MAP TABLE (3 entries TOTAL entries 3)

OXID | RXID | HOST FCID | TARG FCID | VSAN | Index
-----+-----+-----+-----+-----+-----
0xc308|0xffff|0x00690400|0x00620426|0x0005|0x00003364
0xc320|0xffff|0x00690400|0x00620426|0x0005|0x00003365
0xc338|0xffff|0x00690400|0x00620426|0x0005|0x00003366

```

Example 28-17 Displays Information About Tapes for which Exchanges are Tape Accelerated at the Host End FCIP Link

```

switch# show fcip host-tape-session 1

HOST TAPE SESSIONS (1 entries TOTAL entries 1)

Host Tape Session #1
  FCID 0xef0001, VSAN 1, LUN 0x0002
  Outstanding Exchanges 0, Outstanding Writes 0
  Target End Buffering 0 Bytes, Auto Max Writes 1
  Flags 0x0, FSM state Non TA Mode
  First index 0xffffffff7, Last index 0xffffffff7
  Current index=0xfffffffffe, Els Oxid 0xffff7, Seq-Id 0x0000
  Hosts 1
    FCID 0x20300

```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 28-18 Displays Information About Tapes for which Exchanges are Tape Accelerated at the Target End FCIP Link

```
switch# show fcip target-tape-session 2

TARGET TAPE SESSIONS (1 entries  TOTAL entries 1)

Target Tape Session #1
  FCID 0xef0001, VSAN 2, LUN 0x0002
  Outstanding Exchanges 0, Outstanding Writes 0
  Estimated IO Time 0x0
  Flags 0x0, Timer Flags 0x0
  First index 0xffffffff7, Last index 0xffffffff7
  Current index=0xffffffffe, Els Oxid 0xffff7, Seq-Id 0x0000
  Hosts 1
    FCID 0x20300
```

Example 28-19 Displays the FCIP Interface Summary of Counters for a Specified Interface

```
switch# show interface fcip 10
fcip10 is up
  Hardware is GigabitEthernet
  Port WWN is 20:d0:00:0c:85:90:3e:80
  Peer port WWN is 20:d4:00:0c:85:90:3e:80
  Admin port mode is auto, trunk mode is on
  Port mode is E, FCID is 0x720000
  Port vsan is 91
  Speed is 1 Gbps
  Using Profile id 91  (interface GigabitEthernet4/1)
  Peer Information
    Peer Internet address is 3.3.3.2 and port is 3225
  Write acceleration mode is off
  Tape acceleration mode is off
  Tape Accelerator flow control buffer size is 256 KBytes
  IP Compression is disabled
  Special Frame is disabled
  Maximum number of TCP connections is 2
  Time Stamp is disabled
  QOS control code point is 0
  QOS data code point is 0
  B-port mode disabled
  TCP Connection Information
    50529025 Active TCP connections
      Local 0.0.0.7:6, Remote 0.0.0.200:0
    0 host table full 0 target entries in use
    211419104 Attempts for active connections, 1500 close of connections
  TCP Parameters
    Path MTU 124160 bytes
    Current retransmission timeout is 124160 ms
    Round trip time: Smoothed 127829 ms, Variance: 14336
    Advertized window: Current: 0 KB, Maximum: 14 KB, Scale: 14336
    Peer receive window: Current: 0 KB, Maximum: 0 KB, Scale: 51200
    Congestion window: Current: 14 KB, Slow start threshold: 49344 KB
    Current Send Buffer Size: 206463 KB, Requested Send Buffer Size: 429496728
3 KB
    CWM Burst Size: 49344 KB
    5 minutes input rate 491913172779207224 bits/sec, 61489146597400903 bytes/sec, 0 frames/sec
    5 minutes output rate 491913175298921320 bits/sec, 61489146912365165 bytes/sec, 14316551 frames/sec
    5702 frames input, 482288 bytes
    5697 Class F frames input, 481736 bytes
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```

    5 Class 2/3 frames input, 552 bytes
    0 Reass frames
    0 Error frames timestamp error 0
5704 frames output, 482868 bytes
5698 Class F frames output, 482216 bytes
    6 Class 2/3 frames output, 652 bytes
    0 Error frames

```

Example 28-20 Displays Detailed FCIP Interface Standard Counter Information

```

switch# show interface fcip 4 counters
fcip4
    TCP Connection Information
...
    5 minutes input rate 207518944 bits/sec, 25939868 bytes/sec, 12471 frames/sec
    5 minutes output rate 205340328 bits/sec, 25667541 bytes/sec, 12340 frames/sec
    2239902537 frames input, 4658960377152 bytes
        18484 Class F frames input, 1558712 bytes
        2239884053 Class 2/3 frames input, 4658958818440 bytes
        0 Reass frames
        0 Error frames timestamp error 0
    2215051484 frames output, 4607270186816 bytes
        18484 Class F frames output, 1558616 bytes
        2215033000 Class 2/3 frames output, 4607268628200 bytes
        0 Error frames

```

The txbytes is the amount of data before compression. After compression, the compressed txbytes bytes are transmitted with compression and the uncompressed txbytes bytes are transmitted without compression. A packet may be transmitted without compression, if it becomes bigger after compression (see [Example 28-21](#)).

Example 28-21 Displays Detailed FCIP Interface Compression Information, if Enabled

```

switch# show interface fcip 4 counters
fcip4
    TCP Connection Information
...
    IP compression statistics
        208752 rxbytes, 208752 rxbytes compressed
        5143584 txbytes
            0 txbytes compressed, 5143584 txbytes non-compressed
            1.00 tx compression ratio

```

Example 28-22 Displays Detailed FCIP Interface Write Acceleration Counter Information, if Enabled

```

switch# show interface fcip 4 counters
fcip4
    TCP Connection Information
...
    Write Accelerator statistics
        6091 packets in      5994 packets out
        0 frames dropped    0 CRC errors
        0 rejected due to table full
        0 ABTS sent         0 ABTS received
        0 tunnel synchronization errors
        37 writes recd      37 XFER_RDY sent (host)
        0 XFER_RDY rcvd (target)
        37 XFER_RDY rcvd (host)
        0 XFER_RDY not proxied due to flow control (host)
        0 bytes queued for sending
        0 estimated bytes queued on the other side for sending

```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
0 times TCP flow ctrl(target)
0 bytes current TCP flow ctrl(target)
```

Example 28-23 Displays Detailed FCIP Interface Tape Acceleration Counter Information, if Enabled

```
switch# show interface fcip 4 counters
fcip4
  TCP Connection Information
...
  Tape Accelerator statistics
    1 (host) tape sessions      0 (target) tape sessions
    37 writes recd      33 STATUS proxied (host)
    37 write good STATUS rcvd (host)
    0 write good STATUS rcvd (target)
    0 write bad STATUS rcvd (host)
    0 write bad STATUS rcvd (target)
    4 writes not TAed      8 queued flow ctrl (host)
    0 recovery REC sent  Got 0 ACCs  0 Rejects (host)
    0 ABTS sent  Got 0 ACCs (host)
    0 REC Accs  0 REC Rjts      14 REC fwded (host)
    0 SRR Accs  0 SRR Rjts      0 SRR fwded(host)
    0 XferRdy retries      0 Status retries (host)
    0 recovery REC sent  Got 0 ACCs  0 Rejects (target)
    0 recovery SRR sent  Got 0 ACCs (target)
    0 ABTS sent  Got 0 ACCs (target)
    0 tmf cmds rcvd (host)
    0 tmf cmds rcvd (target)
```

Example 28-24 Displays the Compression Engine Statistics for the MPS-14/2 Module

```
switch# show ips stats hw-comp all
HW Compression Statistics for port GigabitEthernet3/1
  Compression stats
    0 input bytes,  0 output compressed bytes
    0 input pkts,   0 output compressed pkts
  Decompression stats
    0 input compressed bytes, 0 output bytes
    0 input compressed pkts,  0 output pkts
  Passthru stats
    0 input bytes, 0 output bytes
    0 input pkts,  0 output pkts
  Miscellaneous stats
    32 min input pktlen,  32 max input pktlen
    28 min output pktlen, 28 max output pktlen
    0 len mismatch,      0 incomplete processing
    0 invalid result,    0 invalid session drop
    0 comp expanded
HW Compression Statistics for port GigabitEthernet3/2
  Compression stats
    0 input bytes,  0 output compressed bytes
    0 input pkts,   0 output compressed pkts
  Decompression stats
    0 input compressed bytes, 0 output bytes
    0 input compressed pkts,  0 output pkts
  Passthru stats
    0 input bytes, 0 output bytes
    0 input pkts,  0 output pkts
  Miscellaneous stats
    32 min input pktlen,  32 max input pktlen
    28 min output pktlen, 28 max output pktlen
```


Send documentation comments to mdsfeedback-doc@cisco.com.

```
0 len mismatch,      0 incomplete processing
0 invalid result,    0 invalid session drop
0 comp expanded
```

Example 28-25 Displays Brief FCIP Interface Counter Information

```
switch# show interface fcip 3 counters brief
```

```
-----
Interface           Input (rate is 5 min avg)      Output (rate is 5 min avg)
-----
                    Rate      Total                          Rate      Total
                    Mbits/s  Frames                      Mbits/s  Frames
-----
fcip3                9        0                          9        0
-----
```

Example 28-26 Displays the FCIP Interface Description

```
switch# show interface fcip 51 description
```

```
FCIP51
  Sample FCIP interface
```

Example 28-27 Displays FCIP Profiles

```
switch# show fcip profile
```

```
-----
ProfileId           Ipaddr           TcpPort
-----
1                   10.10.100.150    3225
2                   10.10.100.150    3226
40                  40.1.1.2         3225
100                 100.1.1.2        3225
200                 200.1.1.2        3225
-----
```

Example 28-28 Displays the Specified FCIP Profile Information

```
switch# show fcip profile 7
```

```
FCIP Profile 7
  Internet Address is 47.1.1.2 (interface GigabitEthernet4/7)
  Listen Port is 3225
  TCP parameters
    SACK is disabled
    PMTU discovery is enabled, reset timeout is 3600 sec
    Keep alive is 60 sec
    Minimum retransmission timeout is 300 ms
    Maximum number of re-transmissions is 4
    Send buffer size is 0 KB
    Maximum allowed bandwidth is 1000000 kbps
    Minimum available bandwidth is 15000 kbps
    Estimated round trip time is 1000 usec
```

Send documentation comments to mdsfeedback-doc@cisco.com.

FCIP High Availability

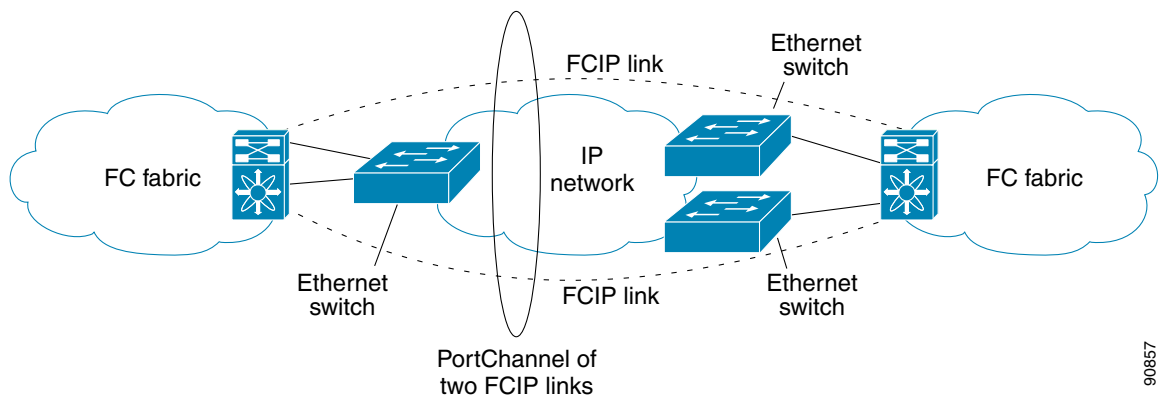
The following high availability solutions are available for FCIP configurations:

- [Fibre Channel PortChannels, page 28-48](#)
- [FSPF, page 28-49](#)
- [VRRP, page 28-49](#)
- [Ethernet PortChannels, page 28-50](#)

Fibre Channel PortChannels

[Figure 28-15](#) provides an example of a PortChannel-based load balancing configuration. To perform this configuration, you need two IP addresses on each SAN island. This solution addresses link failures.

Figure 28-15 PortChannel Based Load Balancing



The following characteristics set Fibre Channel PortChannel solutions apart from other solutions:

- The entire bundle is one logical (E)ISL link.
- All FCIP links in the PortChannel should be across the same two switches.
- The Fibre Channel traffic is load balanced across the FCIP links in the PortChannel.


90857

Send documentation comments to mdsfeedback-doc@cisco.com.

FSPF

Figure 28-16 displays a FSPF-based load balancing configuration example. This configuration requires two IP addresses on each SAN island, and addresses IP and FCIP link failures.

Figure 28-16 FSPF-Based Load Balancing




The following characteristics set FSPF solutions apart from other solutions:

- Each FCIP link is a separate (E)ISL.
- The FCIP links can connect to different switches across two SAN islands.
- The Fibre Channel traffic is load balanced across the FCIP link.

VRRP

Figure 28-17 displays a VRRP-based high availability FCIP configuration example. This configuration requires at least two physical Gigabit Ethernet ports connected to the Ethernet switch on the island where you need to implement high availability using VRRP.

Figure 28-17 VRRP-Based High Availability



The following characteristics set VRRP solutions apart from other solutions:


- If the active VRRP port fails, the standby VRRP port takes over the VRRP IP address.
- When the VRRP switchover happens, the FCIP link automatically disconnects and reconnects.
- This configuration has only one FCIP (E)ISL link.

Send documentation comments to mdsfeedback-doc@cisco.com.

Ethernet PortChannels

Figure 28-18 displays an Ethernet PortChannel-based high availability FCIP example. This solution addresses the problem caused by individual Gigabit Ethernet link failures.

Figure 28-18 Ethernet PortChannel-Based High Availability



The following characteristics set Ethernet PortChannel solutions apart from other solutions:

- The Gigabit Ethernet link level redundancy ensures a transparent failover if one of the Gigabit Ethernet links fails.
- Two Gigabit Ethernet ports in one Ethernet PortChannel appear like one logical Gigabit Ethernet link.
- The FCIP link stays up during the failover.


Ethernet PortChannels and Fibre Channel PortChannels

Ethernet PortChannels offer link redundancy between the Cisco MDS 9000 Family switch's Gigabit Ethernet ports and the connecting ethernet switch. On the other hand, Fibre Channel PortChannels offer (E)ISL link redundancy between Fibre Channel switches. FCIP is an (E)ISL link and is only applicable for a Fibre Channel PortChannel. Beneath the FCIP level, an FCIP link can run on top of an Ethernet PortChannel or just on one Gigabit Ethernet port. This link is totally transparent to the Fibre Channel layer.

An Ethernet PortChannel restriction only allows two contiguous IPS ports, such as ports 1–2 or 3–4, to be combined in one Ethernet Portchannel (see the [“Configuring Gigabit Ethernet High Availability” section on page 28-15](#)). This restriction only applies to Ethernet PortChannels. The Fibre Channel PortChannel (to which FCIP link can be a part of), does not have a restriction on which (E)ISL links can be combined in a Fibre Channel PortChannel as long as it passes the compatibility check (see the [“Compatibility Check” section on page 14-10](#)). The maximum number of Fibre Channel ports that can be put into a Fibre Channel PortChannel is 16 (see [Figure 28-19](#)).

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 28-19 PortChannels at the Fibre Channel and Ethernet Levels

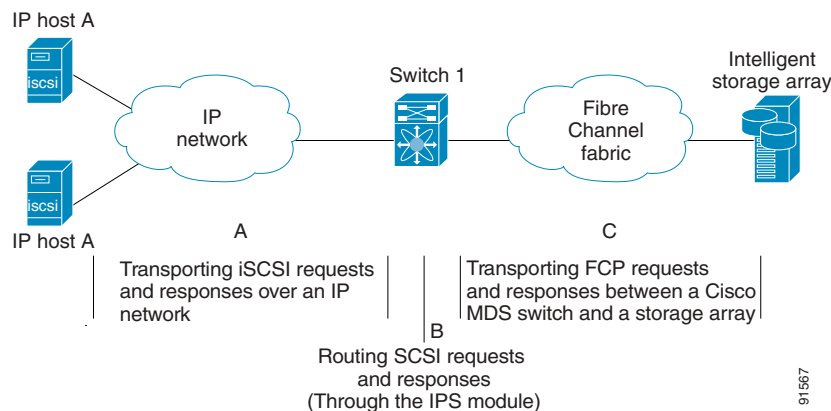


To configure Fibre Channel PortChannels, see [Chapter 14, “Configuring PortChannels.”](#) To configure Ethernet PortChannels, see the [“About Ethernet PortChannel Aggregation”](#) section on page 28-17.

Configuring iSCSI

The iSCSI feature consists of routing iSCSI requests and responses between iSCSI hosts in an IP network and Fibre Channel storage devices in the Fibre Channel SAN that are accessible from any Fibre Channel interface of the Cisco MDS 9000 Family switch (see [Figure 28-20](#)).

Figure 28-20 Transporting iSCSI Requests and Responses for Transparent iSCSI Routing



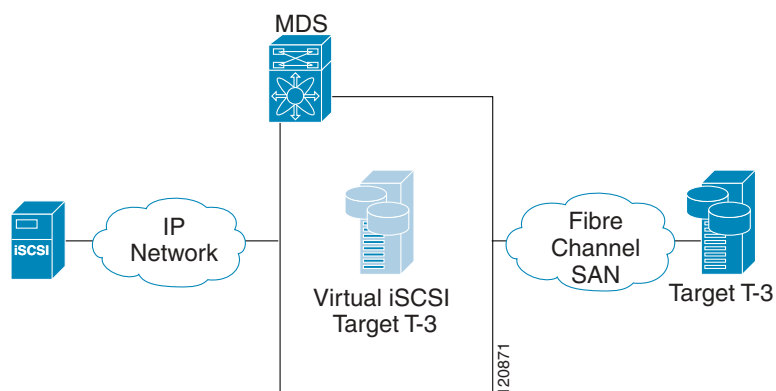
Each iSCSI host that requires access to storage through the IPS module or MPS-14/2 module needs to have a compatible iSCSI driver installed. (The Cisco.com website at <http://www.cisco.com/cgi-bin/tablebuild.pl/sn5420-scsi> provides a list of compatible drivers). Using the iSCSI protocol, the iSCSI driver allows an iSCSI host to transport SCSI requests and responses over an IP network. From the host operating system perspective, the iSCSI driver appears to be a SCSI transport driver similar to a Fibre Channel driver in the host.

Send documentation comments to mdsfeedback-doc@cisco.com.

The IPS module or MPS-14/2 module provides transparent SCSI routing. IP hosts using the iSCSI protocol can transparently access targets on the Fibre Channel network. [Figure 28-20](#) provides an example of a typical configuration of iSCSI hosts connected to an IPS module or MPS-14/2 module via the IP network access Fibre Channel storage on the Fibre Channel SAN.

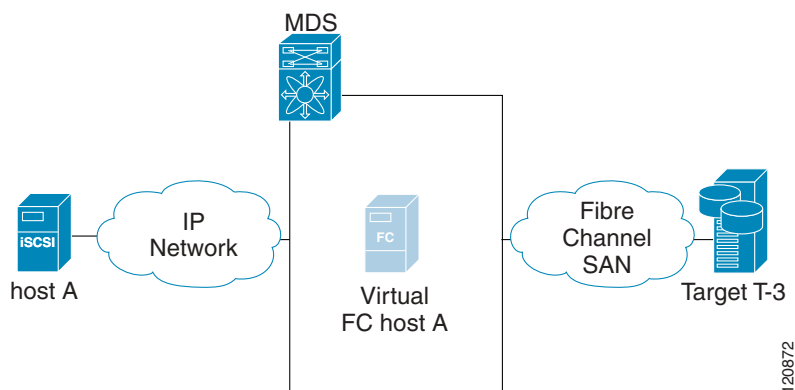
The IPS module or MPS-14/2 module create a separate iSCSI SAN view and Fibre Channel SAN view. For the iSCSI SAN view, the IPS module or MPS-14/2 module create iSCSI virtual targets and then maps them to physical Fibre Channel targets available in the Fibre Channel SAN. They present the Fibre Channel targets to IP hosts as if the physical iSCSI targets were attached to the IP network (see [Figure 28-21](#)).

Figure 28-21 iSCSI SAN View—iSCSI virtual targets



For the Fibre Channel SAN view, the IPS module or MPS-14/2 module presents iSCSI hosts as a virtual Fibre Channel host. The storage devices communicate with the virtual Fibre Channel host similar to communications performed with real Fibre Channel hosts (see [Figure 28-22](#)).


Figure 28-22 Fibre Channel SAN View—iSCSI Host as an HBA



The IPS modules or MPS-14/2 modules transparently map the command between the iSCSI virtual target and the virtual Fibre Channel host (see [Figure 28-23](#)).

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 28-23 iSCSI to FCP (Fibre Channel) Routing



Routing SCSI from the IP host to the Fibre Channel storage device consists of the following main actions:

- The iSCSI requests and responses are transported over an IP network between the hosts and the IPS module or MPS-14/2 module.
- The SCSI requests and responses are routed between the hosts on an IP network and the Fibre Channel storage device (converting iSCSI to FCP and vice versa). The IPS module or MPS-14/2 module performs this conversion and routing.
- The FCP requests or responses are transported between the IPS module or MPS-14/2 module and the Fibre Channel storage devices.



Note

FCP (the Fibre Channel equivalent of iSCSI) carries SCSI commands over a Fibre Channel SAN. Refer to the IETF standards for IP storage at <http://www.ietf.org> for information on the iSCSI protocol.

Enabling iSCSI

To use the iSCSI feature, you must explicitly enable iSCSI on the required switches in the fabric. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family.

To enable iSCSI on any participating switch, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# iscsi enable | Enables iSCSI on that switch. |
| | switch(config)# no iscsi enable | Disables (default) iSCSI on that switch. |



Caution

When you disable this feature, all related configurations are automatically discarded.

Creating iSCSI Interfaces

Each physical GigabitEthernet interface on an IPS module or MPS-14/2 module can be used to translate and route iSCSI requests to Fibre Channel targets and responses in the opposite direction. To enable this capability, the corresponding iSCSI interface must be in an enabled state.

Send documentation comments to mdsfeedback-doc@cisco.com.

To enable iSCSI interfaces, follow these steps:

Step 1 Enable the required Gigabit Ethernet interface.

```
switch# config terminal
switch(config)# interface gigabitethernet 2/1
switch(config-if)# no shutdown
```

Step 2 Create the required iSCSI interface and enable the interface.

```
switch(config-if)# exit
switch(config)# interface iscsi 2/1
switch(config-if)# no shutdown
```

Presenting Fibre Channel Targets as iSCSI Targets

The IPS module or MPS-14/2 module presents physical Fibre Channel targets as iSCSI virtual targets allowing them to be accessed by iSCSI hosts. It does this in one of two ways:

- Dynamic mapping—automatically maps all the Fibre Channel target devices/ports as iSCSI devices. Use this mapping to create automatic iSCSI target names.
- Static mapping—Manually create iSCSI target devices and map them to the whole Fibre Channel target port or a subset of Fibre Channel LUNs. With this mapping, you must specify unique iSCSI target names.

Static mapping should be used when iSCSI hosts should be restricted to subsets of LUs in the Fibre Channel targets and/or iSCSI access control is needed (see the [“iSCSI Access Control” section on page 28-67](#)). Also, static mapping allows configuration of transparent failover if the LUs of the Fibre Channel targets are reachable by redundant Fibre Channel ports (see the [“Transparent Target Failover” section on page 28-86](#)).



Note

The IPS module or MPS-14/2 module does not import Fibre Channel targets to iSCSI by default. Either dynamic or static mapping must be configured before the IPS module or MPS-14/2 module makes Fibre Channel targets available to iSCSI initiators.

Dynamic Mapping

When you configure dynamic mapping the IPS module or MPS-14/2 module imports all Fibre Channel targets to the iSCSI domain and maps each physical Fibre Channel target port as one iSCSI target. That is, all LU accessible through the physical storage target port are available as iSCSI LUs with the same LU number (LUN) as in the physical Fibre Channel target port.

The iSCSI target node name is created automatically using the iSCSI qualified name (IQN) format. The iSCSI qualified name is restricted to a maximum name length of 223 alphanumeric characters and a minimum length of 16 characters.

Send documentation comments to mdsfeedback-doc@cisco.com.

The IPS module or MPS-14/2 module creates an IQN formatted iSCSI target node name using the following conventions because the name must be unique in the SAN:

- IPS Gigabit Ethernet ports that are not part of a VRRP group or PortChannel use this format:
iqn.1987-05.com.cisco:05.<mgmt-ip-address>.<slot#>-<port#>-<sub-intf#>.<Target-pWWN>
- IPS ports that are part of a VRRP group use this format:
iqn.1987-05.com.cisco:05.vrrp-<vrrp-ID#>-<vrrp-IP-addr>.<Target-pWWN>
- Ports that are part of a PortChannel use this format:
iqn.1987-02.com.cisco:02.<mgmt-ip-address>.pc-<port-ch-sub-intf#>.<Target-pWWN>




Note

If you have configured a switch name, then the switch name is used instead of the management IP address. If you have not configured a switch name, the management IP address is used.

With this convention, each IPS port in a Cisco MDS 9000 Family switch creates a unique iSCSI target node name for the same Fibre Channel target port in the SAN.

For example, if an iSCSI target was created for a Fibre Channel target port with pWWN 31:00:11:22:33:44:55:66 and that pWWN contains LUN 0, LUN 1, and LUN 2, those LUNs would become available to an IP host via the iSCSI target node name iqn.1987-05.com.cisco:05.<MDS_switch_management_IP_address>.01-01.3100112233445566 (see [Figure 28-24](#)).

Figure 28-24 Dynamic Target Mapping



Note

Each iSCSI initiator may not have access to all targets depending on the configured access control mechanisms (see [“iSCSI Access Control”](#) section on page 28-67).

To enable dynamic importing of Fibre Channel targets into iSCSI, follow these steps:

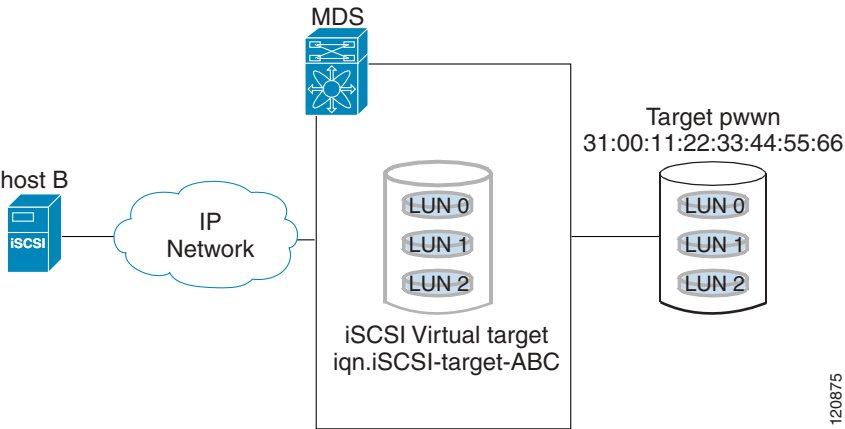
| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# conf ig terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# iscsi import target fc | IPS modules and MPS-14/2 modules dynamically import all Fibre Channel targets in the Fibre Channel SAN into the IP network. |

[Send documentation comments to mdsfeedback-doc@cisco.com.](mailto:mdsfeedback-doc@cisco.com)

Static Mapping

You can manually (statically) create an iSCSI target by assigning a user-defined unique iSCSI node name to it. The iSCSI qualified name is restricted to a minimum length of 16 characters and a maximum of 223 characters. A statically mapped iSCSI target can either map the whole Fibre Channel target port (all LUNs in the target port mapped to the iSCSI target), or it can contain one or more LUs from a Fibre Channel target port (see Figure 28-20).

Figure 28-25 Statically Mapped iSCSI Targets



To create a static iSCSI virtual target for the entire Fibre Channel target port, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# iscsi virtual-target name iqn.iSCSI-target-abc switch(config-iscsi-tgt)# | Creates the iSCSI target name iqn.iSCSI-target-abc. |
| Step 3 | switch(config-iscsi-tgt)# pwwn 31:00:11:22:33:44:55:66 | Maps the Fibre Channel target port to the iSCSI virtual target. One iSCSI target cannot contain more than one Fibre Channel target port. Do not specify the LUN if you wish to map the whole Fibre Channel target to an iSCSI target. All Fibre Channel target LUNs are exposed to iSCSI. |
| | switch(config-iscsi-tgt)# pwwn 31:00:11:22:33:44:55:66 fc-lun 1 iscsi-lun 1 | Maps a virtual target using LUN mapping options. Use this LUN option to map different Fibre Channel LUNs to different iSCSI virtual targets. |

Tip An iSCSI target cannot contain more than one Fibre Channel target port. If you have already mapped the whole Fibre Channel target port, you cannot use the LUN mapping option.

Send documentation comments to mdsfeedback-doc@cisco.com.

**Note**

See the “iSCSI ACL Based Access Control” section on page 28-69 for more information on controlling access to statically-mapped targets.

Advertising Static iSCSI Targets

You can limit the Gigabit Ethernet interfaces via which static iSCSI targets are advertised. By default iSCSI targets are advertised on all Gigabit Ethernet interfaces, subinterfaces, PortChannel interfaces, and PortChannel subinterfaces.

To configure a specific interface that should advertise the iSCSI virtual target, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | <code>switch(config-iscsi-tgt)# advertise interface GigabitEthernet 2/5</code> | Advertises the virtual target only on the specified interface. By default, it is advertised on all interfaces in all IPS modules or MPS-14/2 modules. Note To advertise the virtual target on multiple interfaces, issue the command for each interface. |
| | <code>switch(config-iscsi-tgt)# no advertise interface GigabitEthernet 2/5</code> | Removes this interface from the list of interfaces from which this target is advertised. |


iSCSI Virtual Target Configuration Examples

This section provides three examples of iSCSI virtual target configurations.

Example 1

This example assigns the whole Fibre Channel target as a iSCSI virtual target. All LUNs that are part of the Fibre Channel target are available as part of the iSCSI target (see [Figure 28-26](#)).

Figure 28-26 Assigning iSCSI Node Names




```
iscsi virtual-target name iqn.1987-02.com.cisco.target-1
pWWN 28:00:01:02:03:04:05:06
```

Example 2

This example maps a subset of LUNs of a Fibre Channel target to three iSCSI virtual targets. Each iSCSI target only has one LUN (see [Figure 28-27](#)).

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 28-27 Mapping LUNs to a iSCSI Node Name




```
iscsi virtual-target name iqn.1987-02.com.cisco.target-1
  pWWN 28:00:01:02:03:04:05:06 fc-lun 0 iscsi-lun 0
iscsi virtual-target name iqn.1987-02.com.cisco.target-2
  pWWN 28:00:01:02:03:04:05:06 fc-lun 1 iscsi-lun 0
iscsi virtual-target name iqn.1987-02.com.cisco.target-3
  pWWN 28:00:01:02:03:04:05:06 fc-lun 2 iscsi-lun 0
```

Example 3

This example maps three subsets of Fibre Channel LUN targets to three iSCSI virtual targets. Two iSCSI targets have one LUN and the third iSCSI target has two LUNs (see [Figure 28-27](#)).

Figure 28-28 Mapping LUNs to Multiple iSCSI Node Names



```
iscsi virtual-target name iqn.1987-02.com.cisco.target-1
  pWWN 28:00:01:02:03:04:05:06 fc-lun 0 iscsi-lun 0
iscsi virtual-target name iqn.1987-02.com.cisco.target-2
  pWWN 28:00:01:02:03:04:05:06 fc-lun 1 iscsi-lun 0
iscsi virtual-target name iqn.1987-02.com.cisco.target-3
  pWWN 28:00:01:02:03:04:05:06 fc-lun 2 iscsi-lun 0
  pWWN 28:00:01:02:03:04:05:06 fc-lun 3 iscsi-lun 1
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Presenting iSCSI Hosts as Virtual Fibre Channel Hosts

The IPS module or MPS-14/2 module connects to the Fibre Channel storage devices on behalf of the iSCSI host to send commands and transfer data to and from the storage devices. These modules use a virtual Fibre Channel N port to access the Fibre Channel storage devices on behalf of the iSCSI host. iSCSI hosts are identified by either iSCSI qualified name (IQN) or IP address.

Initiator Identification

iSCSI hosts can be identified by the IPS module or MPS-14/2 module using the following:

- iSCSI qualified name (IQN)

An iSCSI initiator is identified based on the iSCSI node name it provides in the iSCSI login. This mode can be useful if an iSCSI host has multiple IP address and you want to provide the same service independent of the IP address used by the host. An initiator with multiple IP addresses (multiple network interface cards—NICs) has one virtual N port on each IPS port to which it logs in.

- IP address

An iSCSI initiator is identified based on the IP address of the iSCSI host. This mode is useful if an iSCSI host has multiple IP addresses and you want to provide different service based on the IP address used by the host. It is also easier to get the IP address of a host compared to getting the iSCSI node name. A virtual N port is created for each IP address it uses to log in to iSCSI targets. If the host using one IP address logs in to multiple IPS ports, each IPS port will create one virtual N port for that IP address.

You can configure the iSCSI initiator identification mode on each IPS port and all the iSCSI hosts terminating on the IPS port will be identified according to that configuration. The default mode is to identify the initiator by name.

To specify the initiator identification mode, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# interface iscsi 4/1 switch(config-if)# | Selects the iSCSI interface on the switch that identifies all the initiators. |
| Step 3 | switch(config-if)# switchport initiator id ip-address | Identifies the iSCSI initiator based on the IP address. |
| | switch(config-if)# switchport initiator id name | Identifies the iSCSI initiator based on the initiator node name. This is the default behavior. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Initiator Presentation Modes

Two modes are available to present iSCSI hosts in the Fibre Channel fabric: transparent initiator mode and proxy initiator mode.


- In transparent initiator mode, each iSCSI host is presented as one virtual Fibre Channel host. The benefit of transparent mode is it allows a finer-level of Fibre Channel access control configuration (similar to managing a “real” Fibre Channel host). Because of the one-to-one mapping from iSCSI to Fibre Channel, each host can have different zoning or LUN access control on the Fibre Channel storage device.
- In proxy-initiator mode, there is only one virtual Fibre Channel host per one IPS port and all iSCSI hosts use that to access Fibre Channel targets. In a scenario where the Fibre Channel storage device requires explicit LUN access control for every host, the static configuration for each iSCSI initiator can be overwhelming. In such case, using the proxy-initiator mode simplifies the configuration.

Transparent Initiator Mode

Each iSCSI host is presented as one virtual Fibre Channel host (that is, one Fibre Channel N port). The benefit of transparent mode is it allows a finer-level of Fibre Channel access control configuration. Because of the one-to-one mapping from iSCSI to Fibre Channel, each host can have different zoning or LUN access control on the Fibre Channel storage device.

When an iSCSI host connects to the IPS module or MPS-14/2 module, a virtual host N port (HBA port) is created for the host (see [Figure 28-28](#)). Every Fibre Channel N port requires a unique Node WWN and Port WWN.

Figure 28-29 Virtual Host HBA Port



After the virtual N port is created with the WWNs, a Fabric Login (FLOGI) is done via the virtual iSCSI interface of the IPS port. After the FLOGI is completed, the virtual N port is on-line in the Fibre Channel SAN and virtual N port is registered in the Fibre Channel Name Server. The IPS module or MPS-14/2 module registers the following entries in the Fibre Channel Name Server:

- IP address of the iSCSI host in the IP-address field on the Name Server
- IQN of the iSCSI host in the symbolic-node-name field of the Name Server

Send documentation comments to mdsfeedback-doc@cisco.com.

- SCSI_FCP in the FC-4 type field of the name server
- Initiator flag in the FC-4 feature of the name server
- Vendor-specific iSCSI GW flag in the FC-4 type field to identify the N-port device as a iSCSI gateway device in the NS.

When all the iSCSI sessions from the iSCSI host are terminated, the IPS modules or MPS-14/2 modules perform an explicit Fabric logout (FLOGO) to remove the virtual N-port device from the Fibre Channel SAN (this indirectly de-registers the device from the Fibre Channel Name Server).

For every iSCSI session from the host to the iSCSI virtual target there is a corresponding Fibre Channel session to the real Fibre Channel target. In [Figure 28-29](#), there are three iSCSI hosts and all three of them connect to the same Fibre Channel target. There is one Fibre Channel session from each of the three virtual Fibre Channel hosts to the target.

iSCSI Initiator Idle Timeout

iSCSI initiator idle timeout specifies the time for which the virtual Fibre Channel N port is kept idle after the initiator logs out from its last iSCSI session. The default value for this timer is 300 seconds. This is useful to avoid N ports logging in to and logging off of the Fibre Channel SAN as transient failure occurs in the IP network. This helps reduce unnecessary RSCNs being generated in the Fibre Channel SAN.

To configure the initiator idle timeout, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# conf t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# iscsi initiator idle-timeout 10 | Configures the iSCSI initiators to have an idle timeout value of 10 seconds. |

WWN Assignment for iSCSI Initiators

An iSCSI host is mapped to an N port's WWNs by one of the following mechanisms:

- Dynamic mapping (default)
- Static mapping

Dynamic Mapping

With dynamic mapping, an iSCSI host is mapped to a dynamically generated port WWN (pWWN) and node WWN (nWWN). Each time the iSCSI host connects it might be mapped to a different WWN. Use this option if no access control is required on Fibre Channel target device (because the target device access control is usually configured using the host WWN).

The WWNs are allocated from the MDS switch's WWN pool. The WWN mapping to the iSCSI host is maintained as long as the iSCSI host has at least one iSCSI session to the IPS port. When all iSCSI sessions from the host are terminated and the IPS module or MPS-14/2 module performs an FLOGO for the virtual N port of the host, the WWNs are released back to the switch's Fibre Channel WWN pool. These addresses are then available for assignment to other iSCSI hosts requiring access to the Fibre Channel Fabric.

Dynamic mapping is the default mode of operation.

Send documentation comments to mdsfeedback-doc@cisco.com.

Static Mapping

With static mapping, an iSCSI host is mapped to a specific pWWN and nWWN. This mapping is maintained in persistent storage and each time the iSCSI host connects, the same WWN mapping is used. This mode is required if you use access control on the target device.

You can implement static mapping in one of two ways:

- User assignment—You can specify your own unique WWN by providing them during the configuration process.
- System assignment—You can request that the switch should provide a WWN from the switch's Fibre Channel WWN pool and keep the mapping in its configuration.



Tip

We recommend using the **system-assign** option. If you manually assign a WWN, you must ensure its uniqueness (see the “[Configuring World Wide Names](#)” section on page 39-18). You should not use any previously-assigned WWNs.

To configure static mapping (using the **name** option) for an iSCSI initiator, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# iscsi initiator name iqn.1987-02.com.cisco.initiator switch(config-iscsi-init)# | Configures an iSCSI initiator using the iSCSI name of the initiator node. The maximum name length is restricted to 223 alphanumeric characters. The minimum length is 16. |
| | switch(config)# no iscsi initiator name iqn.1987-02.com.cisco.initiator | Deletes the configured iSCSI initiator. |

To configure static mapping (using the **ip-address** option) for an iSCSI initiator, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# iscsi initiator ip-address 10.50.0.0 switch(config-iscsi-init)# | Configures an iSCSI initiator using the IP address of the initiator node. |
| | switch(config)# no iscsi initiator ip-address 10.50.0.0 | Deletes the configured iSCSI initiator. |

To assign the WWN for an iSCSI initiator, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch(config-iscsi-init)# static nwwn system-assign | Uses the switch's WWN pool to allocate the nWWN for this iSCSI initiator and keeps it persistent. |
| | switch(config-iscsi-init)# static nwwn 20:00:00:05:30:00:59:11 | Assigns the user provided WWN as nWWN for the iSCSI initiator. You can only specify one nWWN for each iSCSI node. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|---|---|
| Step 2 | <code>switch(config-iscsi-init)# static pwwn system-assign 2</code> | Uses the switch's WWN pool to allocate two pWWNs for this iSCSI initiator and keeps it persistent. The range is from 1 to 64. |
| | <code>switch(config-iscsi-init)# static pwwn 21:00:00:20:37:73:3b:20</code> | Assigns the user provided WWN as pWWN for the iSCSI initiator. |



Note

If the system-assign option is used to configure WWNs for an iSCSI initiator, when the configuration is saved to an ASCII file the system-assigned WWNs are also saved. Subsequently if you perform a **write erase**, you must manually delete the WWN configuration from the ASCII file. Failing to do so can cause duplicate WWN assignments if the ASCII configuration file is reapplied on the switch.

Making the Dynamic Initiator WWN Mapping Static

After a dynamic initiator has already logged in, you may decide to permanently keep the automatically assigned nWWN/pWWN mapping so this initiator uses the same mapping the next time it logs in.

To permanently keep the automatically assigned nWWN/pWWN mapping, follow these steps:

| | Command | Purpose |
|---------|--|--|
| Step 1 | <code>switch# config terminal</code> <code>switch(config)#</code> | Enters configuration mode. |
| Step 2s | <code>switch(config)# iscsi save-initiator name iqn.1987-02.com.cisco.initiator</code> | Saves the nWWN and pWWNs that have automatically been assigned to the iSCSI initiator whose name is specified. |
| | <code>switch(config)# iscsi save-initiator ip-address 10.10.100.11</code> | Saves the nWWN and pWWNs that have automatically been assigned to the iSCSI initiator whose IP address is specified. |
| | <code>switch(config)# iscsi save-initiator</code> | Saves the nWWN and pWWNs that have automatically been assigned to all the initiators. |
| Step 3 | <code>switch(config)# exit</code> <code>switch#</code> | Returns to EXEC mode. |
| Step 4 | <code>switch# copy running-config startup-config</code> | Saves the nWWN/pWWN mapping configuration across system reboots. |


Proxy-Initiator Mode

In a scenario where the Fibre Channel storage device requires explicit LUN access control for every host using the transparent initiator mode (presenting one iSCSI host as one Fibre Channel host) means every iSCSI host has to be configured statically. This can mean several configuration tasks for each iSCSI host. In this case, using the proxy-initiator mode simplifies the configuration.

In this mode, only one virtual host N port (HBA port) is created per IPS port. All the iSCSI hosts connecting to that IPS port will be multiplexed using the same virtual host N port (see [Figure 28-30](#)). This mode simplifies the task of statically binding WWNs. LUN mapping and assignment on the Fibre Channel storage array must be configured to allow access from the proxy virtual N port's pWWN for all LUNs used by each iSCSI initiator which connects through this IPS port. The LUN is then assigned to each iSCSI initiator by configuring iSCSI virtual targets (see the [“Static Mapping” section on page 28-56](#)) with LUN mapping and iSCSI access control (see the [“iSCSI Access Control” section on page 28-67](#)).

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 28-30 Multiplexing IPS Ports



Proxy initiator mode can be configured on a per IPS port basis, in which case only iSCSI initiators terminating on that IPS port will be in this mode.

When a IPS port is configured in proxy-initiator mode, Fabric Login (FLOGI) is done via the virtual iSCSI interface of the IPS port. After the FLOGI is completed, the proxy-initiator virtual N port is on-line in the Fibre Channel Fabric and virtual N port is registered in the Fibre Channel Name Server. The IPS module or MPS-14/2 module registers the following entries in the Fibre Channel Name Server:

- iSCSI interface name iSCSI slot /port is registered in the symbolic-node-name field Name Server
- SCSI_FCP in the FC-4 type field of the name server
- Initiator flag in the FC-4 feature of the name server
- Vendor specific flag (iscsi-gw) in the FC-4 type field to identify the N-port device as a iSCSI gateway device in the name server

Similar to transparent initiator mode, user can provide a pWWN and nWWN or request a system assigned WWN for the proxy initiator N port.

To configure the proxy initiator, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# interface iscsi 4/1 switch(config-if)# | Selects the iSCSI interface on the switch that initiators will connect to. |
| Step 3 | switch(config-if)# switchport proxy-initiator | Configures the proxy initiator mode with system-assignment nWWN and pWWN. |
| | switch(config-if)# no switchport proxy-initiator | Disables the proxy initiator mode. |
| Step 4 | switch(config-if)# switchport proxy-initiator nWWN 11:11:11:11:11:11:11:11 pwwn 22:22:22:22:22:22:22:22 | (Optional) Configures the proxy initiator mode using the specified WWNs. |
| | switch(config-if)# no switchport proxy-initiator nWWN 11:11:11:11:11:11:11:11 pwwn 22:22:22:22:22:22:22:22 | Disables the proxy initiator mode. |

Send documentation comments to mdsfeedback-doc@cisco.com.


Note

When an interface is in proxy initiator mode, you can only configure Fibre Channel access control (zoning) based on the iSCSI interface's proxy N port attributes—the WWN pairs or the FCID. You cannot configure zoning using iSCSI attributes such as IP address or IQN of the iSCSI initiator. To enforce initiator-based access control, use iSCSI based access control (see the [“iSCSI Access Control” section on page 28-67](#)).

VSAN Membership for iSCSI

Similar to Fibre Channel devices, iSCSI devices have two mechanisms by which VSAN membership can be defined.

- iSCSI host—VSAN membership to iSCSI host (This method takes precedent over the iSCSI interface.)
- iSCSI interface—VSAN membership to iSCSI interface (All iSCSI hosts connecting to this iSCSI interface inherit the interface VSAN membership if the host is not configured in any VSAN by the iSCSI host method.)

VSAN Membership for iSCSI Hosts

Individual iSCSI hosts can be configured to be in a specific VSAN (similar to the DPVM feature for Fibre Channel, see [Chapter 11, “Creating Dynamic VSANs”](#)). The specified VSAN overrides the iSCSI interface VSAN membership.

To assign VSAN membership for iSCSI hosts, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# iscsi initiator name iqn.1987-02.com.cisco.initiator switch(config-iscsi-init)# | Configures an iSCSI initiator. |
| Step 3 | switch(config-iscsi-init)# vsan 3 | Assigns the iSCSI initiator node to a specified VSAN. |
| | switch(config-iscsi-init)# no vsan 5 | Note You can assign this host to one or more VSANs. Removes the iSCSI node from the specified VSAN. |


Note

When an initiator is configured in any other VSAN (other than VSAN 1), for example VSAN 2, the initiator is automatically removed from VSAN 1. If you also want it to be present in VSAN 1, you must explicitly configure the initiator in VSAN 1.

Send documentation comments to mdsfeedback-doc@cisco.com.

VSAN Membership for iSCSI Interfaces

VSAN membership can be configured for an iSCSI interface, called the port VSAN. All the iSCSI devices that connect to this interface automatically become members of this VSAN, if it is not explicitly configured in a VSAN. In other words, the port VSAN of an iSCSI interface is the default VSAN for all dynamic iSCSI initiators. The default port VSAN of an iSCSI interface is VSAN 1.



Tip

This feature was introduced in Cisco SAN-OS Release 1.3(1). If you downgrade to an earlier release, be sure to delete any assigned VSAN and to issue the **no iscsi interface vsan-membership** command before performing the downgrade procedure.

To change the default port VSAN, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# iscsi interface vsan-membership | Enables you to configure VSAN membership for iSCSI interfaces. |
| Step 3 | switch(config)# vsan database switch(config-vsan-db)# | Configures the database for a VSAN. Application specific VSAN parameters cannot be configured from this prompt. |
| Step 4 | switch(config-vsan-db)# vsan 2 interface iscsi 2/1 | Assigns the membership of the iscsi 2/1 interface to the specified VSAN (VSAN 2). |
| | switch(config-vsan-db)# no vsan 2 interface iscsi 2/1 | Reverts to using the default VSAN as the port VSAN of the iSCSI interface. |


Example of VSAN membership for iSCSI devices

Figure 28-31 provides an example of VSAN membership for iSCSI devices:

- iSCSI interface 1/1 is member of VSAN Y
- iSCSI initiator host-A has explicit VSAN membership to VSAN X
- Three iSCSI initiators host-A, host-B, Host-C connect to iSCSI interface 1/1

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 28-31 VSAN Membership for iSCSI Interfaces



Host A's virtual Fibre Channel N port will be added to VSAN X because of explicit membership for the initiator. The virtual host-B and host-C N ports do not have any explicit membership configuration so will inherit the iSCSI interface VSAN membership and be part of VSAN Y.

Advanced VSAN membership for iSCSI hosts

An iSCSI host can be a member of multiple VSANs. In this case multiple virtual Fibre Channel host are created, one in each VSAN in which the iSCSI host is a member. This configuration is useful when certain resources such as Fibre Channel tape devices need to be shared among different VSANs.

iSCSI Access Control

Two mechanisms of access control are available for iSCSI devices.

- Fibre Channel based zoning access control
- iSCSI ACL-based access control

Depending on the initiator mode used to present the iSCSI hosts in the Fibre Channel Fabric, either or both the access control mechanisms can be used.

Fibre Channel Zoning Based Access Control

SAN-OS VSAN and zoning concepts has been extended to cover both Fibre Channel devices and iSCSI devices. Zoning is the standard access control mechanism for Fibre Channel devices which is applied within the context of a VSAN. Fibre Channel zoning has been extended to support iSCSI devices and their extension has the advantage of having a uniform, flexible access control mechanism across the whole SAN.

Send documentation comments to mdsfeedback-doc@cisco.com.

Common mechanisms of identifying members in Fibre Channel Zone are the following (see [Chapter 15, “Configuring and Managing Zones”](#) for details of Fibre Channel Zoning):

- Fibre Channel device pWWN
- Interface and Switch WWN. Device connecting via that interface is within the zone:

In the case of iSCSI, behind an iSCSI interface multiple iSCSI devices may be connected. Interface based zoning may not be useful because all the iSCSI devices behind the interface will automatically be within the same zone.

In transparent initiator mode (where one Fibre Channel virtual N port is created for each iSCSI host as described in the [“Transparent Initiator Mode”](#) section on page 28-60) if an iSCSI host has static WWN mapping then the standard Fibre Channel device pWWN-based zoning membership mechanism can be used.

Zoning membership mechanism has been enhanced to add iSCSI devices to zones based on the following:

- IP address/mask (IP subnet)
- Symbolic-node-name (IQN)

For iSCSI hosts that do not have a static WWN mapping, the feature allows the IP address or iSCSI node name to be specified as zone members. Note that iSCSI hosts that have static WWN mapping can also use these features. IP address based zone membership allows multiple devices to be specified in one command by providing the subnet mask.



Note

In proxy initiator mode, all iSCSI devices connecting to an IPS port gain access to the Fibre Channel fabric via a single virtual Fibre Channel N port. Thus, zoning based on the iSCSI node name or IP address will not have any effect. If zoning based on pWWN is used then all iSCSI devices connecting to that IPS port will be put in the same zone. To implement individual initiator access control in proxy initiator mode, configure a iSCSI ACL on the virtual target (see the [“iSCSI ACL Based Access Control”](#) section on page 28-69).

To add an iSCSI initiator to the zone database, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# zone name iSCSIzone vsan 1 switch(config-zone) | Creates a zone name for the iSCSI devices in the IPS module or MPS-14/2 module to be included. |
| Step 3 | switch(config-zone)# member symbolic-nodename iqn.1987-02.com.cisco.initiator1 | Assigns an iSCSI node name-based membership into a zone. |
| | switch(config-zone)# no member symbolic-nodename iqn.1987-02.com.cisco.init1 | Deletes the specified device from a zone. |
| | switch(config-zone)# member ip-address 10.50.1.1 | Assigns an iSCSI IP address-based membership into a zone. |
| | switch(config-zone)# no member ip-address 10.50.1.1 | Deletes the identified device from a zone. |
| | switch(config-zone)# member pwwn 20:00:00:05:30:00:59:11 | Assigns an iSCSI port WWN-based membership into a zone. |
| | switch(config-zone)# no member pwwn 20:00:00:05:30:00:59:11 | Deletes the device identified by the port WWN from a zone. |

Send documentation comments to mdsfeedback-doc@cisco.com.

iSCSI ACL Based Access Control

iSCSI based access control is applicable only if static iSCSI virtual targets are created (see the “[Static Mapping](#)” section on page 28-56). For a static iSCSI target, you can configure a list of iSCSI initiators that are allowed to access the targets.

By default, static iSCSI virtual targets are not accessible to any iSCSI host. You must explicitly configure accessibility to allow a iSCSI virtual target to be accessed by all hosts. The initiator access list can contain one or more initiators. The iSCSI initiator can be identified by one of the following mechanisms:

- iSCSI node name
- IP address
- IP subnet



Note

For transparent mode iSCSI initiator, if both Fibre Channel zoning and iSCSI ACLs are used, then for every static iSCSI target that is accessible to the iSCSI host, the initiator’s virtual N port should be in the same Fibre Channel zone as the Fibre Channel target.

To configure access control in iSCSI, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# iscsi virtual-target name iqn.1987-02.com.cisco.initiator switch(config-iscsi-tgt)# | Creates the iSCSI target name iqn.1987-02.com.cisco.initiator. |
| Step 3 | switch(config-iscsi-tgt)# pwwn 26:00:01:02:03:04:05:06 switch(config-iscsi-tgt)# | Maps a virtual target node to a Fibre Channel target. |
| Step 4 | switch(config-iscsi-tgt)# initiator iqn.1987-02.com.cisco.initiator1 permit | Allows the specified iSCSI initiator node to access this virtual target. You can issue this command multiple times to allow multiple initiators. |
| | switch(config-iscsi-tgt)# no initiator iqn.1987-02.com.cisco.initiator1 permit | Prevents the specified initiator node from accessing virtual targets. |
| | switch(config-iscsi-tgt)# initiator ip address 10.50.1.1 permit | Allows the specified IP address to access this virtual target. You can issue this command multiple times to allow multiple initiators. |
| | switch(config-iscsi-tgt)# no initiator ip address 10.50.1.1 permit | Prevents the specified IP address from accessing virtual targets. |
| | switch(config-iscsi-tgt)# initiator ip address 10.50.1.0 255.255.255.0 permit | Allows all initiators in this subnetwork (10.50.1/24) to access this virtual target. |
| | switch(config-iscsi-tgt)# no initiator ip address 10.50.1.0 255.255.255.0 permit | Prevents all initiators in this subnetwork from accessing virtual targets. |
| | switch(config-iscsi-tgt)# all-initiator-permit | Allows all initiator nodes to access this virtual target. |
| | switch(config-iscsi-tgt)# no all-initiator-permit | Prevents any initiator from accessing virtual targets (default). |

Send documentation comments to mdsfeedback-doc@cisco.com.

Enforcing Access Control

IPS modules and MPS-14/2 modules use both iSCSI and Fibre Channel zoning-based access control lists to enforce access control. Access control is enforced both during the iSCSI discovery phase and the iSCSI session creation phase. Access control enforcement is not required during IO phase because the IPS module or MPS-14/2 module is responsible for the routing of iSCSI traffic to Fibre Channel.

- iSCSI discovery phase—When an iSCSI host creates an iSCSI discovery session and queries for all iSCSI targets, the IPS module or MPS-14/2 module returns only the list of iSCSI targets this iSCSI host is allowed to access based on the access control policies discussed in the previous section. The IPS module or MPS-14/2 module does this by querying the Fibre Channel name server for all the devices in the same zone as the initiator in all VSANs. It then filters out the devices that are initiator by looking at the FC4-feature field of the FCNS entry. (If a device does not register as either initiator or target in the FC4-feature field, the IPS module or MPS-14/2 module will advertise it). It then responds to the iSCSI host with the list of targets. Each will have either a static iSCSI target name that you configure or a dynamic iSCSI target name that the IPS module or MPS-14/2 module creates for it (see the [“Dynamic Mapping”](#) section on page 28-54).
- iSCSI session creation—When an IP host initiates an iSCSI session, the IPS module or MPS-14/2 module verifies if the specified iSCSI target (in the session login request) is allowed by both the access control mechanisms described in previous section.

If the iSCSI target is a static mapped target, the IPS module or MPS-14/2 module verifies if the iSCSI host is allowed within the access list of the iSCSI target. If the IP host does not have access, its login is rejected. If the iSCSI host is allowed, it validates if the virtual Fibre Channel N port used by the iSCSI host and the Fibre Channel target mapped to the static iSCSI virtual target are in the same Fibre Channel zone.

If the iSCSI target is an auto-generated iSCSI target, then the IPS module or MPS-14/2 module extracts the WWN of the Fibre Channel target from the iSCSI target name and verifies if the initiator and the Fibre Channel target is in the same Fibre Channel zone or not. If they are, then access is allowed.

The IPS module or MPS-14/2 module uses the Fibre Channel virtual N port of the iSCSI host and does a zone-enforced name server query for the Fibre Channel target WWN. If the FCID is returned by the name server, then the iSCSI session is accepted. Otherwise, the login request is rejected.

iSCSI Session Authentication

The IPS module or MPS-14/2 module supports iSCSI authentication mechanism to authenticate iSCSI hosts that request access to storage. By default, IPS module or MPS-14/2 modules allow CHAP or None authentication of iSCSI initiators. If authentication should always be used, you must configure the switch to allow only CHAP authentication.

For CHAP username or secret validation you can use any method supported and allowed by the Cisco MDS AAA infrastructure (see [Chapter 19, “Configuring Switch Security”](#)). AAA authentication supports RADIUS, TACACS+, or local authentication device.

The **aaa authentication iscsi** command enables aaa authentication for the iSCSI host and specifies the method to use.

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure AAA authentication for an iSCSI user, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# aaa authentication iscsi default group RadServerGrp | Uses RADIUS servers that are added in the group called RadServerGrp for the iSCSI CHAP authentication. |
| | switch(config)# aaa authentication iscsi default group TacServerGrp | Uses TACACS+ servers that are added in the group called TacServerGrp for the iSCSI CHAP authentication. |
| | switch(config)# aaa authentication iscsi default local | Uses the local password database for iSCSI CHAP authentication. |

Authentication Mechanism

You can configure iSCSI CHAP or None authentication at both the global level and at each interface level.

The authentication for a Gigabit Ethernet interface or subinterface overrides the authentication method configured at the global level.

If CHAP authentication should always be used, issue the **iscsi authentication chap** command at either the global level or at a per-interface level. If authentication should not be used at all, issue the **iscsi authentication none** command.

To configure the authentication mechanism for iSCSI, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# iscsi authentication chap | Configures CHAP as the default authentication mechanism globally for the Cisco MDS switch. CHAP authentication is required for all iSCSI sessions. |

To configure the authentication mechanism for iSCSI sessions to a particular interface, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# interface GigabitEthernet 2/1.100 switch(config-if)# | Selects the Gigabit Ethernet interface. |
| Step 3 | switch(config-if)# iscsi authentication none | Specifies that no authentication is required for iSCSI sessions to the selected interface. |

Local Authentication

See the “[Configuring User Accounts](#)” section on page 19-29 to create the local password database. To create new users in the local password database for the iSCSI initiator, the iSCSI keyword is mandatory. Use the **user password password iscsi** command to create a new iSCSI user.

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure iSCSI users for local authentication, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# username iscsiuser password ffsffsfsffs345353554535 iscsi | Configures a user name (iscsiuser) and password (ffsffsfsffs345353554535) in the local database for iSCSI login authentication. |

Restricting iSCSI Initiator Authentication

By default, the iSCSI initiator can use any user name in RADIUS or local database in authenticating itself to the IPS module or MPS-14/2 module (the CHAP user name is independent of the iSCSI initiator name). The IPS module or MPS-14/2 module allows the initiator to login as long as it provides a correct response to the CHAP challenge sent by the switch. This can be a problem if one CHAP user name and password had been compromised.

To restrict an initiator to use a specific user name for CHAP authentication, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# iscsi initiator name iqn.1987-02.com.cisco.init switch(config-iscsi-init)# | Enters the configuration submode for the initiator iqn.1987-02.com.cisco.init. |
| Step 3 | switch(config-iscsi-init)# username user1 | Restricts the initiator iqn.1987-02.com.cisco.init to only authenticate using user1 as its CHAP username. Tip Be sure to define user1 as an iSCSI user in the local AAA database or the RADIUS server. |

Mutual CHAP Authentication

In addition to the IPS module or MPS-14/2 module authentication of the iSCSI initiator, the IPS module or MPS-14/2 module also supports a mechanism for the iSCSI initiator to authenticate the Cisco MDS switch's iSCSI target during the iSCSI login phase. This authentication requires the user to configure a username and password for the switch to present to the iSCSI initiator. The provided password is used to calculate a CHAP response to a CHAP challenge sent to the IPS port by the initiator.

To configure a global iSCSI target username and password to be used by the switch to authenticate itself to an initiator, follow these steps:

| | Command | Purpose |
|--------|--|----------------------------|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|--|---|
| Step 2 | <code>switch(config)# iscsi authentication username testuser password abc123</code> | Configures the switch user account (user1) along with a password (abcd12AAA) specified in clear text (default) for all initiators. The password is limited to 64 characters. |
| | <code>switch(config)# iscsi authentication username user1 password 7 !*asdsfsdfjh!@df</code> | Configures the switch user account (user1) along with the encrypted password (specified by 7) (!@*asdsfsdfjh!@df) for all initiators. |
| | <code>switch(config)# iscsi authentication username user1 password 0 abcd12AAA</code> | Configures the switch user account (user1) along with a password (abcd12AAA) specified in clear text (indicated by 0—default) for all initiators. The password is limited to 64 characters. |
| | <code>switch(config)# no iscsi authentication username testuser</code> | Removes the global configuration for all initiators. |

To configure a per-initiator iSCSI target's user name and password used by the switch to authenticate itself to an initiator, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | <code>switch# config t switch(config)#</code> | Enters configuration mode. |
| Step 2 | <code>switch(config)# iscsi initiator name iqn.1987-02.com.cisco.initiator switch(config-iscsi-init)#</code> | Configures an iSCSI initiator using the iSCSI name of the initiator node. |
| Step 3 | <code>switch(config-iscsi-init)# mutual-chap username testuser password abc123</code> | Configures the switch user account (user1) along with a password (abcd12AAA) specified in clear text (default). The password is limited to 64 characters. |
| | <code>switch(config-iscsi-init)# mutual-chap username user1 password 7 !*asdsfsdfjh!@df</code> | Configures the switch user account (user1) along with the encrypted password (specified by 7) (!@*asdsfsdfjh!@df). |
| | <code>switch(config-iscsi-init)# mutual-chap username user1 password 0 abcd12AAA</code> | Configures the switch user account (user1) along with a password (abcd12AAA) specified in clear text (indicated by 0—default). The password is limited to 64 characters. |
| | <code>switch(config-iscsi-init)# no mutual-chap username testuser</code> | Removes the switch authentication configuration. |

Use the **show running-config** and the **show iscsi global** (see [Example 28-34](#)) commands to display the global configuration. Use the **show running-config** and the **show iscsi initiator configured** (see [Example 28-42](#)) commands to display the initiator specific configuration.

Send documentation comments to mdsfeedback-doc@cisco.com.

iSCSI Interface Advanced Features

Advanced configuration options are available for iSCSI interfaces on a per-IPS port basis. These configurations are similar to the advanced FCIP configurations and are already explained in that section (see the [“Advanced FCIP Interface Configuration” section on page 28-30](#)).

To access these commands from the iSCSI interface, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# interface iscsi 4/1 switch(config-if)# | Selects the iSCSI interface on the switch. |

Cisco MDS switches support the following advanced features for iSCSI interfaces.

iSCSI Listener Port

You can configure the TCP port number for the iSCSI interface which listens for new TCP connections. The default port number is 3260. Once you change the TCP port number, the iSCSI port only accepts TCP connections on the newly configured port.

See the [“Configuring TCP Listener Ports” section on page 28-25](#)

TCP Tuning parameters

You can configure the following TCP parameters.

- The minimum retransmit timeout, keepalive timeout, maximum retransmissions, path MTU, SACK (SACK is enabled by default for iSCSI TCP configurations), window management (The iSCSI defaults are max-bandwidth = 1G, min-available-bandwidth = 70 Mbps, and round-trip-time = 1 ms.), buffer size (default send buffer size for iSCSI is 4096 KB), window congestion (enabled by default and the default burst size is 50 KB.), and maximum delay jitter (enabled by default and the default time is 500 microseconds.).

See the [“Minimum Retransmit Timeout” section on page 28-26](#), [“Keepalive Timeout” section on page 28-26](#), [“Maximum Retransmissions” section on page 28-26](#), [“Path MTUs” section on page 28-27](#), [“Monitoring Congestion” section on page 28-28](#) and [“Estimating Maximum Jitter” section on page 28-29](#).

QoS

To set the QoS values, follow these steps:

| | Command | Purpose |
|--------|------------------------------------|--|
| Step 1 | switch(config-if)# qos 3 | Configure the differentiated services code point (DSCP) value of 3 to be applied to all outgoing IP packets in this iSCSI interface. The valid range for the iSCSI DSCP value is from 0 to 63. |
| Step 2 | switch(config-if)# no qos 5 | Reverts the switch to its factory default (marks all packets with DSCP value 0). |

Send documentation comments to mdsfeedback-doc@cisco.com.

iSCSI Routing Modes

Cisco MDS 9000 Family switches support multiple iSCSI routing modes. Each mode negotiates different operational parameters, has different advantages and disadvantages, and is suitable for different usages.

- **pass-thru** mode

In **pass-thru** mode, the port on the IPS module or MPS 14/2 module converts and forwards read data frames from the Fibre Channel target to the iSCSI host frame-by-frame without buffering. This means that one data-in frame received is immediately sent out as one iSCSI data-in PDU.

In the opposite direction, the port on the IPS module or MPS 14/2 module limits the maximum size of iSCSI write data-out PDU that the iSCSI host can send to the maximum data size that the Fibre Channel target specifies that it can receive. The result is one iSCSI data-out PDU received sent out as one Fibre Channel data frame to the Fibre Channel target.

The absence of buffering in both directions leads to an advantage of lower forwarding latency. However, a small maximum data segment length usually results in lower data transfer performance from the host due to a higher processing overhead by the host system. Another benefit of this mode is iSCSI data digest can be enabled. This helps protect the integrity of iscsi data carried in the PDU over what TCP checksum offers.

- **store-and-forward** mode (default)

In **store-and-forward** mode, the port on the IPS module or MPS 14/2 module assembles all the Fibre Channel data frames of an exchange to build one large iSCSI data-in PDU before forwarding it to the iSCSI client.

In the opposite direction, the port on the IPS module or MPS 14/2 module does not impose a small data segment size on the host so the iSCSI host can send an iSCSI data-out PDU of any size (up to 256 KB). The port then waits until the whole iSCSI data-out PDU is received before it converts, or splits, the PDU, and forwards Fibre Channel frames to the Fibre Channel target.

The advantage of this mode is higher data transfer performance from the host. The disadvantages are higher transfer latency and that the iSCSI data digest (CRC) cannot be used.



Note The **store-and-forward** mode is the default forwarding mode as of Cisco SAN-OS Release 2.0(1b).

- **cut-through** mode

This **cut-through** improves the read operation performance over **store-and-forward** mode. The port on the IPS module or MPS 14/2 module achieves this by forwarding each Fibre Channel data-in frame to the iSCSI host as it is received without waiting for the whole exchange complete. There is no difference for write data-out operations from **store-and-forward** mode.

Figure 28-32 compares the messages exchanged by the iSCSI routing modes.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 28-32 iSCSI Routing Modes

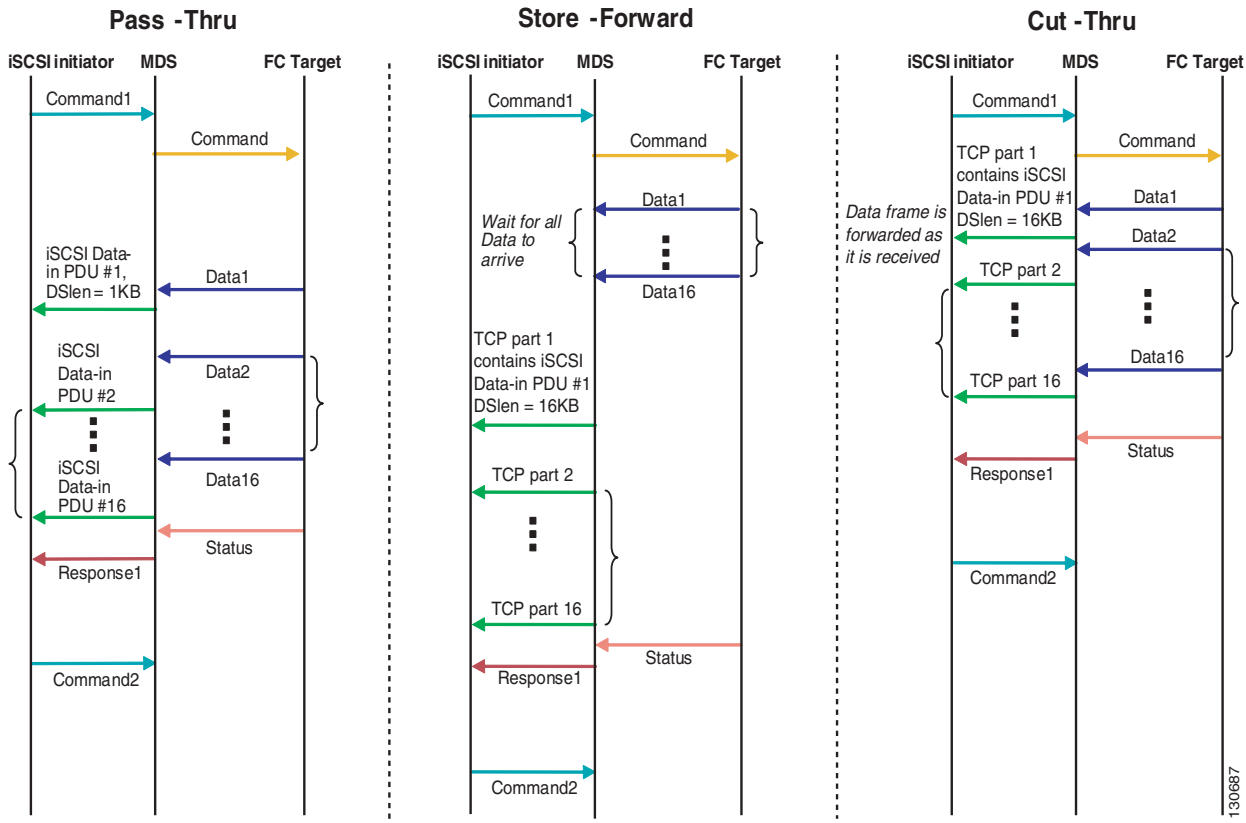


Table 28-2 compares the advantages and disadvantages of the different iSCSI routing modes.

Table 28-2 Comparison of iSCSI Routing Modes

| Mode | Advantages | Disadvantages |
|--------------------------|--|---|
| pass-thru | Low-latency Data digest can be used | Lower data transfer performance |
| store-and-forward | Higher data transfer performance | Data digest cannot be used |
| cut-thru | Improved read performance over store-and-forward | If Fibre Channel target sent read data for different commands interchangeably, data of the first one will be forwarded in cut-thru mode but data of subsequent commands will be buffered and the behavior will be the same as store-and-forward mode. |

To set the iSCSI routing mode, follow these steps:

| | Command | Purpose |
|---------------|--|---|
| Step 1 | <code>switch(config-if)# mode cut-thru</code> | Configures cut-thru mode on the iSCSI interface. |
| Step 2 | <code>switch(config-if)# no mode cut-thru</code> | Reverts store-and-forward mode (default). |

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying iSCSI Information

Use the **show iscsi** command to obtain detailed informations about iSCSI configurations.

Displaying iSCSI Interfaces

Use the **show iscsi interface** command to view the summary, counter, description, and status of the iSCSI interface. Use the output to verify the administrative mode, the interface status, TCP parameters currently used, and brief statistics.

Example 28-29 Displays the iSCSI Interface Information

```
switch# show interface iscsi 4/1
iscsi4/1 is up
  Hardware is GigabitEthernet
  Port WWN is 20:cf:00:0c:85:90:3e:80
  Admin port mode is iSCSI
  Port mode is iSCSI
  Speed is 1 Gbps
  iSCSI initiator is identified by name
  Number of iSCSI session: 0 (discovery session: 0)
  Number of TCP connection: 0
  Configured TCP parameters
    Local Port is 3260
    PMTU discover is enabled, reset timeout is 3600 sec
    Keepalive-timeout is 60 sec
    Minimum-retransmit-time is 300 ms
    Max-retransmissions 4
    Sack is enabled
    QOS code point is 0
    Maximum allowed bandwidth is 1000000 kbps
    Minimum available bandwidth is 70000 kbps
    Estimated round trip time is 1000 usec
    Send buffer size is 4096 KB
    Congestion window monitoring is enabled, burst size is 50 KB
    Configured maximum jitter is 500 us
  Forwarding mode: store-and-forward
  TMF Queueing Mode : disabled
  Proxy Initiator Mode : disabled
  5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  iSCSI statistics
    Input 0 packets, 0 bytes
      Command 0 pdus, Data-out 0 pdus, 0 bytes
    Output 0 packets, 0 bytes
      Response 0 pdus (with sense 0), R2T 0 pdus
      Data-in 0 pdus, 0 bytes
```

Displaying iSCSI Statistics

Use the **show iscsi stats** command to view brief or detailed iSCSI statistics per iSCSI interface. See [Example 28-30](#) and [Example 28-31](#).

[Example 28-30](#) displays iSCSI throughput on an IPS port in both inbound and outbound directions. It also displays the number of different types of iSCSI PDU received and transmitted by this IPS port.

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 28-30 Display brief iSCSI statistics for an iSCSI interface

```
switch# show iscsi stats iscsi 2/1
iscsi2/1
  5 minutes input rate 704 bits/sec, 88 bytes/sec, 1 frames/sec
  5 minutes output rate 704 bits/sec, 88 bytes/sec, 1 frames/sec
  iSCSI statistics
    974756 packets input, 142671620 bytes
      Command 2352 pdus, Data-out 44198 pdus, 92364800 bytes, 0 fragments, unsolicited 0
      bytes
      output 1022920 packets, 143446248 bytes
      Response 2352 pdus (with sense 266), R2T 1804 pdus
      Data-in 90453 pdus, 92458248 bytes
```

Example 28-31 displays detailed iSCSI statistics for an IPS port. Along with the traffic rate and the number of each iSCSI PDU types, it shows the number of FCP frames received and forwarded, the number of iSCSI login attempt, succeed, and failure. It also shows the number of different types of iSCSI PDUs sent and received that are noncritical or occur less frequently, such as NOP in and out (NOP-In and NOP-Out), text request and response (Text-REQ and Text-RESP), and task management request and response (TMF-REQ and TMF-RESP).

Various types of errors and PDU or frame drop occurrences are also counted and displayed. For example, Bad header digest shows the number of iSCSI PDUs received that have a header digest that fails CRC verification. The iSCSI Drop section shows the number of PDUs that were dropped due to reasons such as target down, LUN mapping fail, Data CRC error, or unexpected Immediate or Unsolicited data. These statistics are helpful for debugging purposes when the feature is not working as expected.

The last section, Buffer Stats, gives statistics of the internal IPS packet buffer operation. This section is for debugging purposes only.

Example 28-31 Displays Detailed iSCSI Statistics for the iSCSI Interface

```
switch# show iscsi stats iscsi 2/1 detail
iscsi2/1
  5 minutes input rate 704 bits/sec, 88 bytes/sec, 1 frames/sec
  5 minutes output rate 704 bits/sec, 88 bytes/sec, 1 frames/sec
  iSCSI statistics
    974454 packets input, 142656516 bytes
      Command 2352 pdus, Data-out 44198 pdus, 92364800 bytes, 0 fragments, unsolicited 0
      bytes
      output 1022618 packets, 143431144 bytes
      Response 2352 pdus (with sense 266), R2T 1804 pdus
      Data-in 90453 pdus, 92458248 bytes
  iSCSI Forward:
    Command:2352 PDUs (Rcvd:2352)
    Data-Out (Write):16236 PDUs (Rcvd 44198), 0 fragments, 92364800 bytes, unsolicited 0
  bytes
  FCP Forward:
    Xfer_rdy:1804 (Rcvd:1804)
    Data-In:90453 (Rcvd:90463), 92458248 bytes
    Response:2352 (Rcvd:2362), with sense 266
    TMF Resp:0

  iSCSI Stats:
    Login:attempt:13039, succeed:110, fail:12918, authen fail:0
    Rcvd:NOP-Out:914582, Sent:NOP-In:914582
      NOP-In:0, Sent:NOP-Out:0
      TMF-REQ:0, Sent:TMF-RESP:0
      Text-REQ:18, Sent:Text-RESP:27
      SNACK:0
      Unrecognized Opcode:0, Bad header digest:0
```


Send documentation comments to mdsfeedback-doc@cisco.com.

```

        Command in window but not next:0, exceed wait queue limit:0
        Received PDU in wrong phase:0
        SCSI Busy responses:0
        Immediate data failure::Separation:0
        Unsolicited data failure::Separation:0, Segment:0
            Add header:0
        Sequence ID allocation failure:0
FCP Stats:
    Total:Sent:47654
        Received:96625 (Error:0, Unknown:0)
    Sent:PLOGI:10, Rcvd:PLOGI_ACC:10, PLOGI_RJT:0
        PRLI:10, Rcvd:PRLI_ACC:10, PRLI_RJT:0, Error:0, From initiator:0
        LOGO:4, Rcvd:LOGO_ACC:0, LOGO_RJT:0
        PRLO:4, Rcvd:PRLO_ACC:0, PRLO_RJT:0
        ABTS:0, Rcvd:ABTS_ACC:0
        TMF REQ:0
        Self orig command:10, Rcvd:data:10, resp:10
    Rcvd:PLOGI:156, Sent:PLOGI_ACC:0, PLOGI_RJT:156
        LOGO:0, Sent:LOGO_ACC:0, LOGO_RJT:0
        PRLI:8, Sent:PRLI_ACC:8, PRLI_RJT:0
        PRLO:0, Sent:PRLO_ACC:0, PRLO_RJT:0
        ADISC:0, Sent:ADISC_ACC:0, ADISC_RJT:0
        ABTS:0

iSCSI Drop:
    Command:Target down 0, Task in progress 0, LUN map fail 0
        CmdSeqNo not in window 0, No Exchange ID 0, Reject 0
        No task:0
    Data-Out:0, Data CRC Error:0
    TMF-Req:0, No task:0
    Unsolicited data:0, Immediate command PDU:0
FCP Drop:
    Xfer_rdy:0, Data-In:0, Response:0

Buffer Stats:
    Buffer less than header size:0, Partial:45231, Split:322
    Pullup give new buf:0, Out of contiguous buf:0, Unaligned m_data:0

```

Displaying Proxy Initiator Information

If the proxy initiator feature is enabled in the iSCSI interface, use the **show interface iscsi** command to display configured proxy initiator information (see [Example 28-32](#) and [Example 28-33](#)).

Example 28-32 Displays Proxy Initiator Information for the iSCSI Interface with System-Assigned WWNs

```

switch# show interface iscsi 4/1
iscsi4/1 is up
    Hardware is GigabitEthernet
    Port WWN is 20:c1:00:05:30:00:a7:9e
    Admin port mode is ISCSI
    Port mode is ISCSI
    Speed is 1 Gbps
    iSCSI initiator is identified by name
    Number of iSCSI session: 0, Number of TCP connection: 0
    Configured TCP parameters
        Local Port is 3260
        PMTU discover is enabled, reset timeout is 3600 sec
        Keepalive-timeout is 60 sec
        Minimum-retransmit-time is 300 ms

```

Send documentation comments to mdsfeedback-doc@cisco.com.

```

Max-retransmissions 4
Sack is disabled
QOS code point is 0
Forwarding mode: pass-thru
TMF Queueing Mode : disabled
Proxy Initiator Mode : enabled<-----Proxy initiator is enabled
  nWWN is 28:00:00:05:30:00:a7:a1 (system-assigned)<----System-assigned nWWN
  pWWN is 28:01:00:05:30:00:a7:a1 (system-assigned)<---- System-assigned pWWN
5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
iSCSI statistics
  Input 7 packets, 2912 bytes
    Command 0 pdus, Data-out 0 pdus, 0 bytes
  Output 7 packets, 336 bytes
    Response 0 pdus (with sense 0), R2T 0 pdus
    Data-in 0 pdus, 0 bytes

```

Example 28-33 Displays Proxy Initiator Information for the iSCSI Interface with User-Assigned WWNs

```

switch# show interface iscsi 4/2
iscsi4/2 is up
  Hardware is GigabitEthernet
  Port WWN is 20:c1:00:05:30:00:a7:9e
  Admin port mode is iSCSI
  Port mode is iSCSI
  Speed is 1 Gbps
  iSCSI initiator is identified by name
  Number of iSCSI session: 0, Number of TCP connection: 0
  Configured TCP parameters
    Local Port is 3260
    PMTU discover is enabled, reset timeout is 3600 sec
    Keepalive-timeout is 60 sec
    Minimum-retransmit-time is 300 ms
    Max-retransmissions 4
    Sack is disabled
    QOS code point is 0
  Forwarding mode: pass-thru
  TMF Queueing Mode : disabled
  Proxy Initiator Mode : enabled
    nWWN is 11:11:11:11:11:11:11:11 (manually-configured)<----User-assigned nWWN
    pWWN is 22:22:22:22:22:22:22:22 (manually-configured)<----User-assigned pWWN
5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
iSCSI statistics
  Input 7 packets, 2912 bytes
    Command 0 pdus, Data-out 0 pdus, 0 bytes
  Output 7 packets, 336 bytes
    Response 0 pdus (with sense 0), R2T 0 pdus
    Data-in 0 pdus, 0 bytes

```

Displaying Global iSCSI Information

Use the **show iscsi global** command to view the overall configuration and the iSCSI status. See [Example 28-34](#).

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 28-34 Displays the Current Global iSCSI Configuration and State

```
switch# show iscsi global
iSCSI Global information
  Authentication: CHAP, NONE
  Import FC Target: Enabled
  Initiator idle timeout: 300 seconds
  Number of target node: 0
  Number of portals: 11
  Number of session: 0
  Failed session: 0, Last failed initiator name:
```

Displaying iSCSI Sessions

Use the **show iscsi session** command to view details about the current iSCSI sessions in the switch. Without parameters, this command displays all sessions. The output can be filtered by specifying an initiator, a target, or both.

[Example 28-35](#) displays one iSCSI initiator configured based on the IQN (iqn.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k) and another based on its IP address (10.10.100.199).

Example 28-35 Displays Brief Information of All iSCSI Sessions

```
switch# show iscsi session
Initiator iqn.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k
  Initiator ip addr (s): 10.10.100.116
  Session #1
    Discovery session, ISID 00023d000043, Status active

  Session #2
    Target VT1
    VSAN 1, ISID 00023d000046, Status active, no reservation

  Session #3
    Target VT2
    VSAN 1, ISID 00023d000048, Status active, no reservation

Initiator 10.10.100.199
  Initiator name iqn.1987-05.com.cisco.01.7e3183ae458a94b1cd6bc168cba09d2e
  Session #1
    Target VT2
    VSAN 1, ISID 246700000000, Status active, no reservation

  Session #2
    Target VT1
    VSAN 1, ISID 246b00000000, Status active, no reservation

  Session #3
    Target iqn.1987-05.com.cisco:05.switch.04-01.2100002037a6be32
    VSAN 1, ISID 246e00000000, Status active, no reservation
```

[Example 28-36](#) and [Example 28-37](#) display the iSCSI initiator configured based on its IP address (10.10.100.199).

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 28-36 Displays Brief Information About the Specified iSCSI Session

```
switch# show iscsi session initiator 10.10.100.199 target VT1
Initiator 10.10.100.199
  Initiator name ign.1987-05.com.cisco.01.7e3183ae458a94b1cd6bc168cba09d2e
  Session #1
    Target VT1
      VSAN 1, ISID 246b00000000, Status active, no reservation
```

Example 28-37 Displays Detailed Information About the Specified iSCSI Session

```
switch# show iscsi session initiator 10.10.100.199 target VT1 detail
Initiator 10.10.100.199 (oasis-qa)
  Initiator name ign.1987-05.com.cisco.01.7e3183ae458a94b1cd6bc168cba09d2e
  Session #1 (index 3)
    Target VT1
      VSAN 1, ISID 246b00000000, TSIH 384, Status active, no reservation
      Type Normal, ExpCmdSN 39, MaxCmdSN 54, Barrier 0
      MaxBurstSize 0, MaxConn 0, DataPDUInOrder No
      DataSeqInOrder No, InitialR2T Yes, ImmediateData No
      Registered LUN 0, Mapped LUN 0
    Stats:
      PDU: Command: 38, Response: 38
      Bytes: TX: 8712, RX: 0
    Number of connection: 1
    Connection #1
      Local IP address: 10.10.100.200, Peer IP address: 10.10.100.199
      CID 0, State: LOGGED_IN
      StatSN 62, ExpStatSN 0
      MaxRecvDSLength 1024, our_MaxRecvDSLength 1392
      CSG 3, NSG 3, min_pdu_size 48 (w/ data 48)
      AuthMethod none, HeaderDigest None (len 0), DataDigest None (len 0)
      Version Min: 2, Max: 2
      FC target: Up, Reorder PDU: No, Marker send: No (int 0)
      Received MaxRecvDSLen key: No
```

Displaying iSCSI Initiators

Use the **show iscsi initiator** command to display information about all initiators connected to an iSCSI interface in the switch. The information can be filtered to display only the desired iSCSI initiator by specifying the initiator name. Detailed output of the iscsi initiator can be obtained by specifying the **detail** option. The **iscsi-session** (and optionally **detail**) parameter displays only iSCSI session information. The **fc-session** (and optionally **detail**) parameter displays only FCP session information. The output includes static and dynamic initiators. See [Example 28-38](#) and [Example 28-39](#).

Example 28-38 Displays Information About Connected iSCSI Initiators

```
switch# show iscsi initiator
iSCSI Node name is ign.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k
  Initiator ip addr (s): 10.10.100.116
  iSCSI alias name: AVANTI12-W2K
  Node WWN is 22:01:00:05:30:00:10:e1 (configured)
  Member of vsans: 1, 2, 10
  Number of Virtual n_ports: 1
  Virtual Port WWN is 22:04:00:05:30:00:10:e1 (configured)
    Interface iSCSI 4/1, Portal group tag: 0x180
      VSAN ID 1, FCID 0x6c0202
      VSAN ID 2, FCID 0x6e0000
      VSAN ID 10, FCID 0x790000
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
iSCSI Node name is 10.10.100.199
iSCSI Initiator name: ign.1987-05.com.cisco.01.7e3183ae458a94b1cd6bc168cba09d2e
iSCSI alias name: oasis-qa
Node WWN is 22:03:00:05:30:00:10:e1 (configured)
Member of vsans: 1, 5
Number of Virtual n_ports: 1
Virtual Port WWN is 22:00:00:05:30:00:10:e1 (configured)
  Interface iSCSI 4/1, Portal group tag: 0x180
  VSAN ID 5, FCID 0x640000
  VSAN ID 1, FCID 0x6c0203
```

Example 28-39 Displays Detailed Information About the iSCSI Initiator

```
switch# show iscsi initiator ign.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k detail
iSCSI Node name is ign.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k
  Initiator ip addr (s): 10.10.100.116
  iSCSI alias name: AVANTI12-W2K
  Node WWN is 22:01:00:05:30:00:10:e1 (configured)
  Member of vsans: 1, 2, 10
  Number of Virtual n_ports: 1

Virtual Port WWN is 22:04:00:05:30:00:10:e1 (configured)
  Interface iSCSI 4/1, Portal group tag is 0x180
  VSAN ID 1, FCID 0x6c0202
  1 FC sessions, 1 iSCSI sessions
  iSCSI session details          <-----iSCSI session details
    Target: VT1
    Statistics:
      PDU: Command: 0, Response: 0
      Bytes: TX: 0, RX: 0
      Number of connection: 1
    TCP parameters
      Local 10.10.100.200:3260, Remote 10.10.100.116:4190
      Path MTU: 1500 bytes
      Retransmission timeout: 310 ms
      Round trip time: Smoothed 160 ms, Variance: 38
      Advertized window: Current: 61 KB, Maximum: 62 KB, Scale: 0
      Peer receive window: Current: 63 KB, Maximum: 63 KB, Scale: 0
      Congestion window: Current: 1 KB

  FCP Session details          <-----FCP session details
    Target FCID: 0x6c01e8 (S_ID of this session: 0x6c0202)
    pWWN: 21:00:00:20:37:62:c0:0c, nWWN: 20:00:00:20:37:62:c0:0c
    Session state: CLEANUP
    1 iSCSI sessions share this FC session
    Target: VT1
    Negotiated parameters
      RcvDataFieldSize 1392 our_RcvDataFieldSize 1392
      MaxBurstSize 0, EMPD: FALSE
      Random Relative Offset: FALSE, Sequence-in-order: Yes
    Statistics:
      PDU: Command: 0, Response: 0
```

Use the **show fcns database** (and optionally **detail**) to display the Fibre Channel name server entry for the Fibre Channel N port created for iSCSI initiators in the SAN. See [Example 28-40](#) and [Example 28-41](#).

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 28-40 Displays the FCNS Database Contents

```
switch# show fcns database
VSAN 1:
-----
FCID          TYPE  PWWN                               (VENDOR)          FC4-TYPE:FEATURE
-----
0x020101      N     22:04:00:05:30:00:35:e1 (Cisco)           scsi-fcp:init isc..w <---iSCSI
0x020102      N     22:02:00:05:30:00:35:e1 (Cisco)           scsi-fcp:init isc..w initiator
0x0205d4      NL    21:00:00:04:cf:da:fe:c6 (Seagate)         scsi-fcp:target
0x0205d5      NL    21:00:00:04:cf:e6:e4:4b (Seagate)         scsi-fcp:target
...
Total number of entries = 10

VSAN 2:
-----
FCID          TYPE  PWWN                               (VENDOR)          FC4-TYPE:FEATURE
-----
0xef0001      N     22:02:00:05:30:00:35:e1 (Cisco)           scsi-fcp:init isc..w
Total number of entries = 1

VSAN 3:
-----
FCID          TYPE  PWWN                               (VENDOR)          FC4-TYPE:FEATURE
-----
0xed0001      N     22:02:00:05:30:00:35:e1 (Cisco)           scsi-fcp:init isc..w
Total number of entries = 1
```

Example 28-41 Displays the FCNS Database in Detail

```
switch# show fcns database detail
-----
VSAN:1      FCID:0x020101
-----
port-wwn (vendor)      :22:04:00:05:30:00:35:e1 (Cisco)
node-wwn               :22:03:00:05:30:00:35:e1
class                  :2,3
node-ip-addr           :10.2.2.12                      <--- iSCSI initiator's IP address
ipa                    :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name     :
symbolic-node-name     :iqn.1991-05.com.microsoft:oasis2-dell <--- iSCSI initiator's IQN
port-type              :N
port-ip-addr           :0.0.0.0
fabric-port-wwn        :22:01:00:05:30:00:35:de
hard-addr              :0x000000
-----
VSAN:1      FCID:0x020102
-----
port-wwn (vendor)      :22:02:00:05:30:00:35:e1 (Cisco)
node-wwn               :22:01:00:05:30:00:35:e1
class                  :2,3
node-ip-addr           :10.2.2.11
ipa                    :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name     :
symbolic-node-name     :iqn.1987-05.com.cisco.01.14ac33ba567f986f174723b5f9f2377
port-type              :N
port-ip-addr           :0.0.0.0
fabric-port-wwn        :22:01:00:05:30:00:35:de
hard-addr              :0x000000
...
Total number of entries = 10
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
=====
-----
VSAN:2      FCID:0xef0001
-----
port-wwn (vendor)      :22:02:00:05:30:00:35:e1 (Cisco)
node-wwn               :22:01:00:05:30:00:35:e1
class                 :2,3
node-ip-addr           :10.2.2.11
ipa                   :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name     :
symbolic-node-name     :iqn.1987-05.com.cisco.01.14ac33ba567f986f174723b5f9f2377
port-type              :N
port-ip-addr           :0.0.0.0
fabric-port-wwn        :22:01:00:05:30:00:35:de
hard-addr              :0x000000
Total number of entries = 1
...
```

Use the **show iscsi initiator configured** to display information about all the configured iSCSI initiators. Specifying the name shows information about the desired initiator. See [Example 28-42](#).

Example 28-42 Display Information About Configured Initiators

```
switch# show iscsi initiator configured
iSCSI Node name is iqn.1987-05.com.cisco:02.3021b0f2fda0.avanti12-w2k
  Member of vsans: 1, 2, 10
  Node WWN is 22:01:00:05:30:00:10:e1
  No. of PWWN: 5
    Port WWN is 22:04:00:05:30:00:10:e1
    Port WWN is 22:05:00:05:30:00:10:e1
    Port WWN is 22:06:00:05:30:00:10:e1
    Port WWN is 22:07:00:05:30:00:10:e1
    Port WWN is 22:08:00:05:30:00:10:e1

iSCSI Node name is 10.10.100.199
  Member of vsans: 1, 5
  Node WWN is 22:03:00:05:30:00:10:e1
  No. of PWWN: 4
    Port WWN is 22:00:00:05:30:00:10:e1
    Port WWN is 22:09:00:05:30:00:10:e1
    Port WWN is 22:0a:00:05:30:00:10:e1
    Port WWN is 22:0b:00:05:30:00:10:e1

User Name for Mutual CHAP: testuser
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying iSCSI Virtual Targets

Use the **show iscsi virtual-target** to display information about the Fibre Channel targets exported as iSCSI virtual targets to the iSCSI initiators. The output includes static as well as dynamic targets. See [Example 28-43](#).

Example 28-43 Displays Exported Targets

```
switch# show iscsi virtual-target
target: VT1
  * Port WWN 21:00:00:20:37:62:c0:0c
    Configured node
    all initiator permit is enabled

target: VT2
  Port WWN 21:00:00:04:cf:4c:52:c1
  Configured node
  all initiator permit is disabled
target: iqn.1987-05.com.cisco:05.switch.04-01.2100002037a6be32
  Port WWN 21:00:00:20:37:a6:be:32 , VSAN 1
  Auto-created node
```

Displaying iSCSI User Information

The **show user-account iscsi** command displays all configured iSCSI user names. See [Example 28-44](#).

Example 28-44 Displays iSCSI User Names

```
switch# show user-account iscsi
username:iscsiuser
secret: dsfffsffsffasffsdffg

username:user2
secret: cshadhdhsadadjajdjas
```

iSCSI High Availability

The following high availability features are available for iSCSI configurations:

- [Transparent Target Failover, page 28-86](#)
- [Multiple IPS Ports Connected to the Same IP Network, page 28-90](#)
- [VRRP-Based High Availability, page 28-92](#)
- [Ethernet PortChannel-Based High Availability, page 28-93](#)

Transparent Target Failover

The following high availability configurations are available:

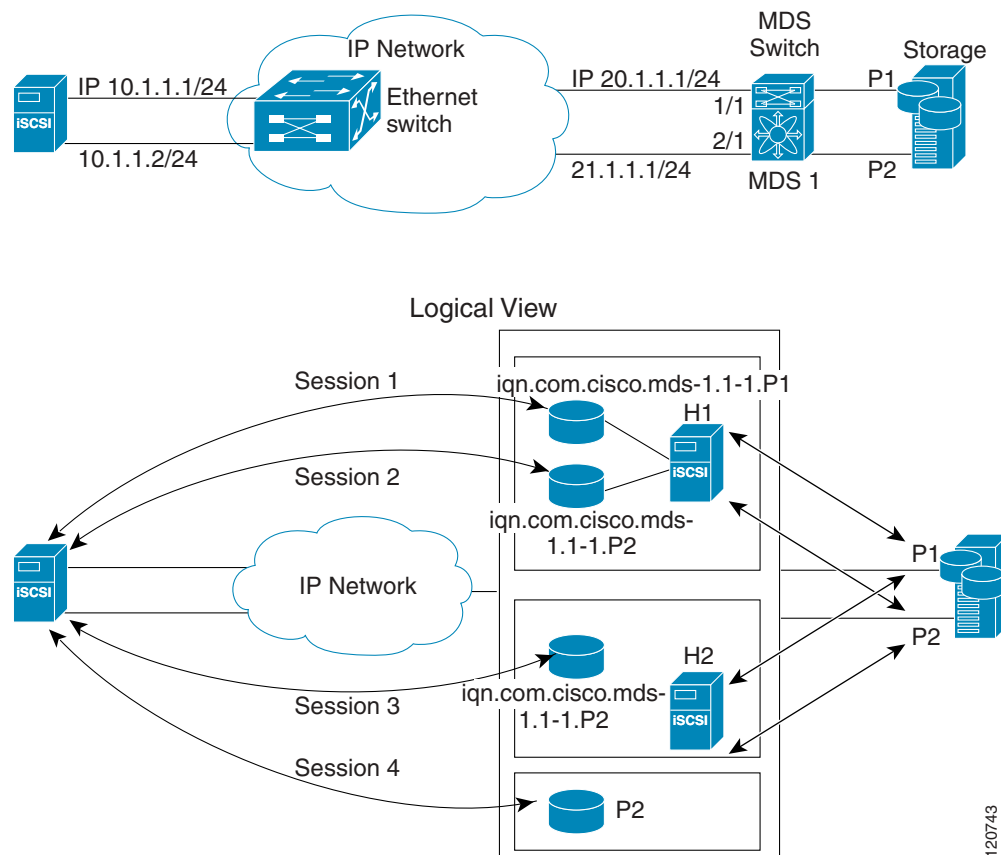
- iSCSI high availability with host running multi-path software
- iSCSI High availability with host not having multi-path software

Send documentation comments to mdsfeedback-doc@cisco.com.

iSCSI High Availability with Host Running Multi-Path Software

Figure 28-33 shows the physical and logical topology for iSCSI HA solution for hosts running multi-path software. In this scenario, the host will have 4 iSCSI sessions. Two iSCSI sessions from each host NIC to the two IPS port.

Figure 28-33 Host Running Multi-Path Software



Each IPS ports is exporting the same two Fibre Channel target ports of the storage but as different iSCSI target names (if you use dynamic iSCSI targets). So the two IPS ports are exporting a total of 4 iSCSI target devices. These four iSCSI targets map the same two ports of the Fibre Channel target.

The iSCSI host uses NIC-1 to connect to IPS port 1 and NIC-2 to connect to IPS port 2. Each IPS port exports two iSCSI targets, so the iSCSI host creates 4 iSCSI sessions.

If the iSCSI host NIC-1 fails (see Figure 28-33 for the physical view), then session 1 and 2 fail but we still have session 3 and 4.

If the IPS port -1 fails, iSCSI host cannot connect to the IPS port, then session 1 and 2 fail. But sessions 3/4 are still available.

If the Storage port P-1 fails, then the IPS ports will terminate session 1 and 3 (put iSCSI virtual target `iqn.com.cisco.mds-5.1-2.p1` and `iqn-com.cisco.mds-5.1-1.p1` in off-line state). But session 2/4 are still available.

In this topology, you have recovery from failure of any of the components. The host multi-path software takes care of load-balancing/fail-over across the different paths to access the storage.

Send documentation comments to mdsfeedback-doc@cisco.com.

iSCSI HA with Host Not Having Any Multi-Path Software


The above topology will not work if the host does not have multi-path software because the host has multiple sessions to the same storage. Without multi-path software the host does not have knowledge the multiple paths are to the same storage.

IP Storage has two additional features to provide HA solution in this scenario.

- IPS ports support VRRP protocol (see the [“Configuring VRRP for Gigabit Ethernet Interfaces” section on page 28-16](#)) to provide failover for IPS ports.
- IPS has transparent Fibre Channel target failover feature for iSCSI static virtual target.

Statically imported iSCSI targets have an additional option to provide a secondary pWWN for the Fibre Channel target. This can be used when the physical Fibre Channel target is configured to have an LU visible across redundant ports. When the active port fails, the secondary port becomes active and the iSCSI session switches to use the new active port (see [Figure 28-34](#)).

Figure 28-34 Static Target Importing Through Two Fibre Channel Ports



In [Figure 28-34](#), you can create a iSCSI virtual target that is mapped to both pWWN1 and pWWN2 to provide redundant access to the Fibre Channel targets.

The failover to secondary port is done transparently by the IPS port without impacting the iSCSI session from the host. All outstanding I/O are terminated with a check condition status when the primary port fails. New I/O received while the failover has not completed will receive a busy status.



Tip

If you use LUN mapping, you can define a different secondary Fibre Channel LUN if the LU number is different.

Enable the optional **revert-primary-port** option to direct the IPS port to switch back to the primary port when the primary port is up again. If this option is disabled (default) and the primary port is up again after a switchover, the old sessions will remain with the secondary port and does not switch back to the primary port. However, any new session will use the primary port. This is the only situation when both the primary and secondary ports are used at the same time.

To create a static iSCSI virtual target, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# iscsi virtual-target name iqn.1987-02.com.cisco.initiator | Creates the iSCSI target name iqn.1987-02.com.cisco.initiator. |

Send documentation comments to mdsfeedback-doc@cisco.com.


| | Command | Purpose |
|--------|---|---|
| Step 3 | switch(config-iscsi-tgt)# pwwn 26:00:01:02:03:04:05:06 secondary-pwwn 26:00:01:02:03:10:11:12 | Configures the primary and secondary ports for this virtual target. |
| Step 4 | switch(config-iscsi-tgt)# revert-primary-port | Configures the session failover redundancy for this virtual-target to switch all sessions back to primary port when the primary port comes back up. |
| | switch(config-iscsi-tgt)# no revert-primary-port | Directs the switch to continue using the secondary port for existing sessions and to use the primary port for new sessions (default) |
| | switch(config-iscsi-tgt)# pwwn 26:00:01:02:03:04:05:06 fc-lun 1 iscsi-lun 0 sec-lun 3 secondary-pwwn 26:00:01:02:03:10:11:12 | Configures the primary and secondary ports for this virtual target with lun mapping and different LU number on the secondary Fibre Channel port. |

Storage Port Failover LUN Trespass

In addition to the high availability of statically imported iSCSI targets, the trespass feature is available (as of Cisco SAN-OS Release 1.3(1)) to enable the move of LUs, on an active port failure, from the active to the passive port of a statically imported iSCSI target.

In physical Fibre Channel targets, which are configured to have LUs visible over two Fibre Channel N ports, when the active port fails, the passive port takes over. Some physical Fibre Channel targets require that the **trespass** option be used to move the LUs from the active port to the passive port. A statically imported iSCSI target's secondary pWWN option and an additional option of enabling the trespass feature is available for a physical Fibre Channel target with redundant ports. When the active port fails, the passive port becomes active, and if the trespass feature is enabled, the Cisco MDS switch sends a request to the target to move the LUs on the new active port. The iSCSI session switches to use the new active port and the moved LUs are accessed over the new active port (see [Figure 28-35](#)).

Figure 28-35 Virtual Target with an Active Primary Port



105219

Send documentation comments to mdsfeedback-doc@cisco.com.

To enable the trespass feature for a static iSCSI virtual target, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# iscsi virtual-target name 1987-02.com.cisco.initiator switch(config-iscsi-tgt)# | Creates the iSCSI target name iqn.1987-02.com.cisco.initiator. |
| Step 3 | switch(config-iscsi-tgt)# pwwn 50:00:00:a1:94:cc secondary-pwwn 50:00:00:a1:97:ac | Maps a virtual target node to a Fibre Channel target and configures a secondary pWWN. |
| Step 4 | switch(config-iscsi-tgt)# trespass | Enables the trespass feature. |
| | switch(config-iscsi-tgt)# no trespass | Disables the trespass feature (default). |

Use the **show iscsi virtual-target** command to verify.

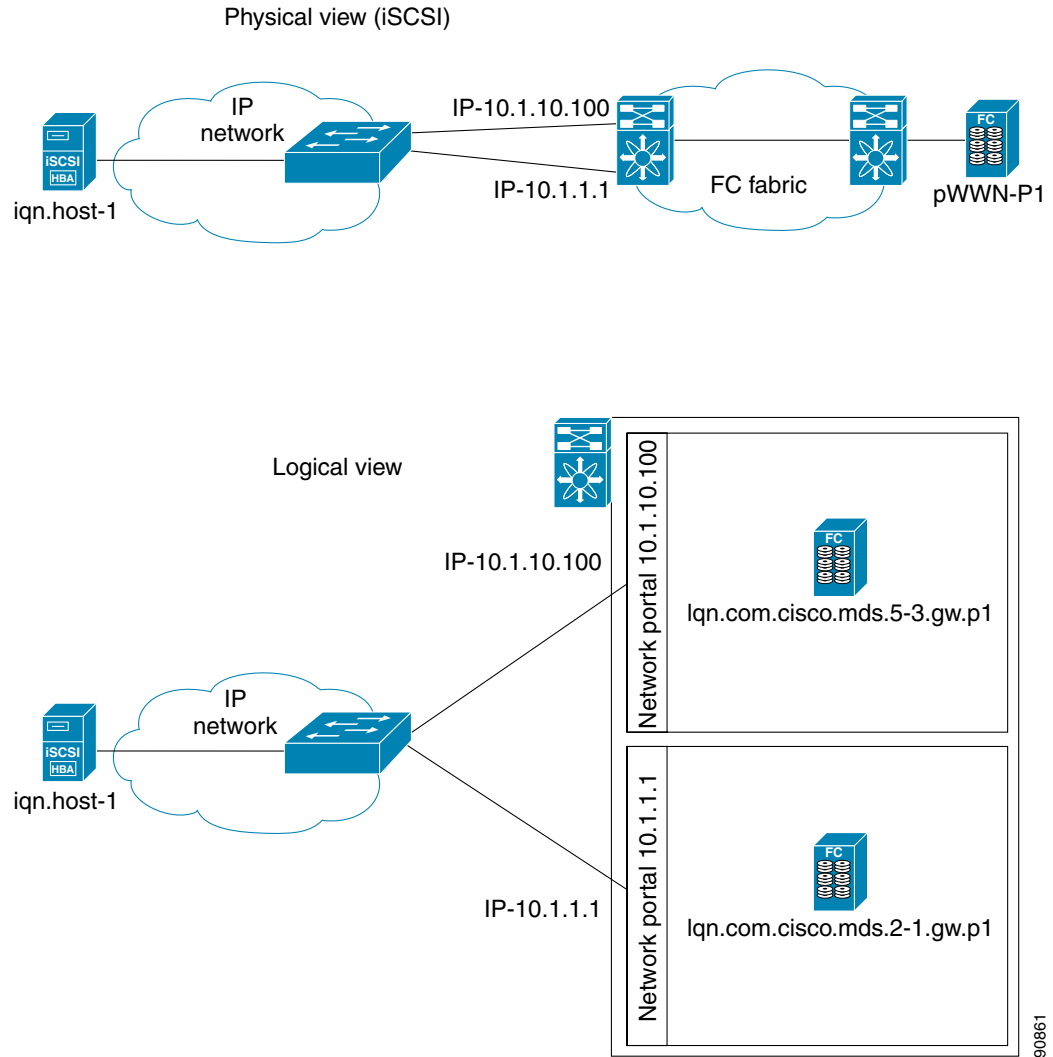
```
switch# show iscsi virtual-target iqn.1987-02.com.cisco.initiator
target: 1987-02.com.cisco.initiator
  Port WWN 10:20:10:00:56:00:70:50
  Configured node
  all initiator permit is disabled
  trespass support is enabled
```

Multiple IPS Ports Connected to the Same IP Network

Figure 28-36 provides an example of a configuration with multiple Gigabit Ethernet interfaces in the same IP network.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 28-36 Multiple Gigabit Ethernet Interfaces in the Same IP Network



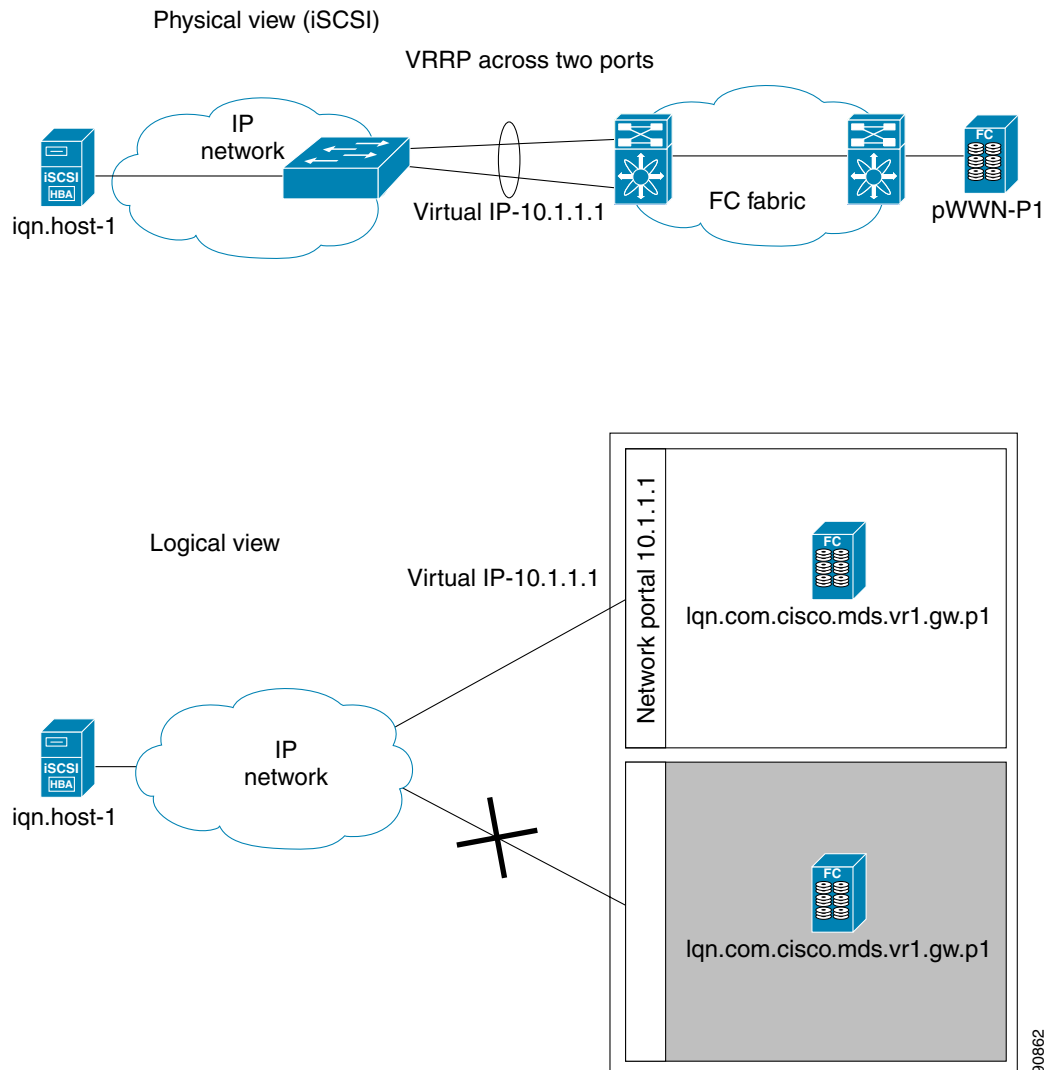
In [Figure 28-36](#), each iSCSI host discovers two iSCSI targets for every physical Fibre Channel target (with different names). The multi-pathing software on the host provides load-balancing over both paths. If one Gigabit Ethernet interface fails, the host multi-pathing software is not affected because it can use the second path.

Send documentation comments to mdsfeedback-doc@cisco.com.

VRRP-Based High Availability

Figure 28-37 provides an example of a VRRP-based high availability iSCSI configuration.

Figure 28-37 VRRP-Based iSCSI High Availability



In Figure 28-37, each iSCSI host discovers one iSCSI target for every physical Fibre Channel target. When the Gigabit Ethernet interface of the VRRP master fails, the iSCSI session is terminated. The host then reconnects to the target and the session comes up because the second Gigabit Ethernet interface has taken over the virtual IP address as the new master.

Send documentation comments to mdsfeedback-doc@cisco.com.

Ethernet PortChannel-Based High Availability




Note

All iSCSI data traffic for one iSCSI link is carried on one TCP connection. Consequently, the aggregated bandwidth is 1 Gbps for that iSCSI link.

Figure 28-38 provides a sample Ethernet PortChannel-based high availability iSCSI configuration.

Figure 28-38 Ethernet PortChannel-Based iSCSI High Availability



In Figure 28-38, each iSCSI host discovers one iSCSI target for every physical Fibre Channel target. The iSCSI session from the iSCSI host to the iSCSI virtual target (on the IPS port) uses one of the two physical interfaces (because an iSCSI session uses one TCP connection). When the Gigabit Ethernet interface fails, the IPS module or MPS-14/2 module and the Ethernet switch transparently forwards all the frames on to the second Gigabit Ethernet interface.

iSCSI Authentication Setup Guidelines and Scenarios

This section provides guidelines on iSCSI authentication possibilities, setup requirements, and sample scenarios. It includes the following authentication setup guidelines:

- [No Authentication, page 28-93](#)
- [CHAP with Local Password Database, page 28-94](#)
- [CHAP with External RADIUS Server, page 28-94](#)



Note

This section does not specify the steps to enter or exit EXEC mode, configuration mode, or any submode. Be sure to verify the prompt before issuing any command.

No Authentication

Set the iSCSI authentication method to **none** to configure a network with no authentication.

```
switch(config)# iscsi authentication none
```

Send documentation comments to mdsfeedback-doc@cisco.com.

CHAP with Local Password Database

To configure authentication using the CHAP option with the local password database, follow these steps:

- Step 1** Set the AAA authentication to use the local password database for iSCSI protocol.

```
switch(config)# aaa authentication iscsi default local
```

- Step 2** Set the iSCSI authentication method to require CHAP for all iSCSI clients.

```
switch(config)# iscsi authentication chap
```

- Step 3** Configure the user names and passwords for iSCSI users.

```
switch(config)# username iscsi-user password abcd iscsi
```



Note If you do not specify the **iscsi** option, the user name is assumed to be a Cisco MDS switch user instead of an iSCSI user.

- Step 4** Verify the global iSCSI authentication setup.

```
switch# show iscsi global
iSCSI Global information Authentication: CHAP <----Verify
  Import FC Target: Disabled
  ...
```

CHAP with External RADIUS Server

To configure authentication using the CHAP option with an external RADIUS server, follow these steps:

- Step 1** Configure the password for the Cisco MDS switch as RADIUS client to the RADIUS server.

```
switch(config)# radius-server key mds-1
```

- Step 2** Configure the RADIUS server IP address.

```
switch(config)# radius-server host 10.1.1.10
```

- Step 3** Configure a server group.

```
switch(config)# aaa group server radius iscsi-radius-group
switch(config-radius)# server 10.1.1.1
```

- Step 4** Set up the authentication verification for the iSCSI protocol to go to the RADIUS server.

```
switch(config)# aaa authentication iscsi default group iscsi-radius-group
```

- Step 5** Set up the iSCSI authentication method to require CHAP for all iSCSI clients.

```
switch(config)# iscsi authentication chap
```

- Step 6** Verify that the global iSCSI authentication set up is CHAP.

```
switch# show iscsi global
iSCSI Global information
  Authentication: CHAP          <----- Verify CHAP
  ....
```


Send documentation comments to mdsfeedback-doc@cisco.com.

Step 7 Verify that the AAA authentication information for iSCSI.

```
switch# show aaa authentication
      default: local
      console: local
      iscsi: group iscsi-radius-group   <----- Group name
      dhchap: local

switch# show radius-server groups
total number of groups:2

following RADIUS server groups are configured:
  group radius:
    server: all configured radius servers
  group iscsi-radius-group:
    server: 10.1.1.1 on auth-port 1812, acct-port 1813

switch# show radius-server
Global RADIUS shared secret:mds-1   <----- Verify secret
....

following RADIUS servers are configured:
  10.1.1.1:   <----- Verify the server IP address
    available for authentication on port:1812
    available for accounting on port:1813
```

To configure an iSCSI RADIUS server, follow these steps:

-
- Step 1** Configure the RADIUS server to allow access from the Cisco MDS switch's management Ethernet IP address.
 - Step 2** Configure the shared secret for the RADIUS server to authenticate the Cisco MDS switch.
 - Step 3** Configure the iSCSI users and passwords on the RADIUS server.
-

iSCSI Transparent Mode Initiator

This scenario assumes the following configuration (see [Figure 28-39](#)):

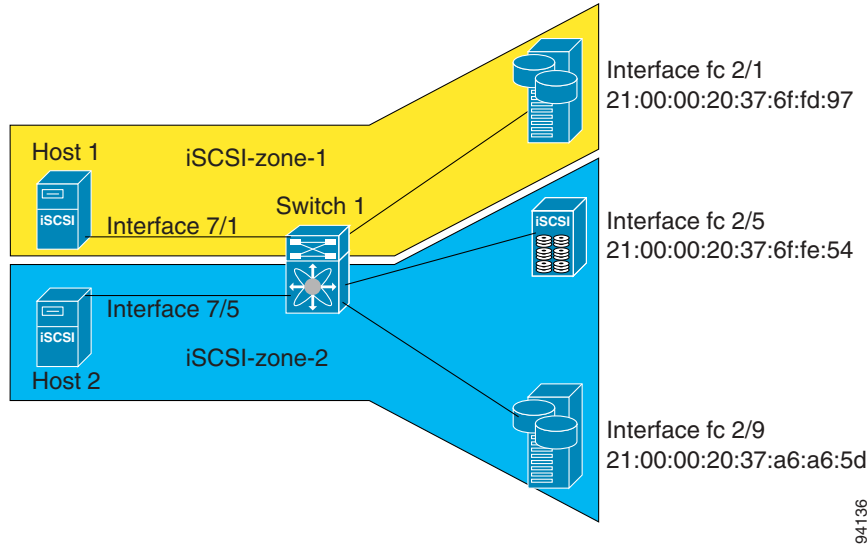
- No LUN mapping or LUN masking or any other access control for hosts on the target device
- No iSCSI login authentication (that is, login authentication set to none)

Topology:

- iSCSI interface 7/1 is configured to identify initiators by IP address
- iSCSI interface 7/5 is configured to identify initiators by node name
- iSCSI initiator host-1 with IP address 10.11.1.10 and name `iqn.1987-05.com.cisco:01.255891611111` connects to IPS port 7/1 is identified using IP address (host 1 = 10.11.1.10).
- The iSCSI initiator host-2 with IP address 10.15.1.10 and node name `iqn.1987-05.com.cisco:01.25589167f74c` connects to IPS port 7/5

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 28-39 iSCSI Scenario 1



To configure scenario 1 (see [Figure 28-39](#)), follow these steps:

- Step 1** Configure null authentication for all iSCSI hosts in Cisco MDS switches.

```
switch(config)# iscsi authentication none
```

- Step 2** Configure iSCSI to dynamically import all Fibre Channel targets into the iSCSI SAN using auto-generated iSCSI target names.

```
switch(config)# iscsi import target fc
```

- Step 3** Configure the Gigabit Ethernet interface in slot 7 port 1 with an IP address and enable the interface.

```
switch(config)# int gigabitethernet 7/1
switch(config-if)# ip address 10.11.1.1 255.255.255.0
switch(config-if)# no shut
```



Note Host 2 is connected to this port.

- Step 4** Configure the iSCSI interface in slot 7 port 1 to identify all dynamic iSCSI initiators by the IP address, and enable the interface.

```
switch(config)# int iscsi 7/1
switch(config-if)# switchport initiator id ip-address
switch(config-if)# no shut
```

- Step 5** Configure the Gigabit Ethernet interface in slot 7 port 5 with the IP address and enable the interface.

```
switch(config)# int gigabitethernet 7/5
switch(config-if)# ip address 10.15.1.1 255.255.255.0
switch(config-if)# no shut
```

- Step 6** Configure the iSCSI interface in slot 7 port 5 to identify all dynamic iSCSI initiators by node name and enable the interface.

```
switch(config)# int iscsi 7/5
switch(config-if)# switchport initiator id name
switch(config-if)# no shut
```

Send documentation comments to mdsfeedback-doc@cisco.com.



Note Host 1 is connected to this port.

Step 7 Verify the available Fibre Channel targets (see [Figure 28-39](#)).

```
switch# show fcns database
```

```
VSAN 1:
```

| FCID | TYPE | PWWN | (VENDOR) | FC4-TYPE:FEATURE |
|-----------------------------|------|-------------------------|-----------|------------------|
| 0x6d0001 | NL | 21:00:00:20:37:6f:fd:97 | (Seagate) | scsi-fcp:target |
| 0x6d0101 | NL | 21:00:00:20:37:6f:fe:54 | (Seagate) | scsi-fcp:target |
| 0x6d0201 | NL | 21:00:00:20:37:a6:a6:5d | (Seagate) | scsi-fcp:target |
| Total number of entries = 3 | | | | |

Step 8 Create a zone named *iscsi-zone-1* with host 1 and one Fibre Channel target in it.



Note Use the IP address of the host in zone membership configuration because the iSCSI interface is configured to identify all hosts based on IP address.

```
switch(config)# zone name iscsi-zone-1 vsan 1
switch(config-zone)# member pwn 21:00:00:20:37:6f:fd:97
switch(config-zone)# member ip-address 10.11.1.10
```

Step 9 Create a zone named *iscsi-zone-2* with host 2 and two Fibre Channel targets in it.



Note Use the symbolic node name of the iSCSI host in zone membership configuration because the iSCSI interface is configured to identify all hosts based on node name.

```
switch(config)# zone name iscsi-zone-2 vsan 1
switch(config-zone)# member pwn 21:00:00:20:37:6f:fe:54
switch(config-zone)# member pwn 21:00:00:20:37:a6:a6:5d
switch(config-zone)# member symbolic-nodename ign.1987-05.com.cisco:01.25589167f74c
```

Step 10 Create a zone set and add the two zones as members.

```
switch(config)# zoneset name zoneset-iscsi vsan 1
switch(config-zoneset)# member iscsi-zone-1
switch(config-zoneset)# member iscsi-zone-2
```

Step 11 Activate the zone set.

```
switch(config)# zoneset activate name zoneset-iscsi vsan 1
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Step 12 Display the active zone set.



Note The iSCSI hosts has not connected so they do not have a FCID yet.

```
switch# show zoneset active
zoneset name zoneset-iscsi vsan 1
  zone name iscsi-zone-1 vsan 1
    * fcid 0x6d0001 [pwwn 21:00:00:20:37:6f:fd:97] <-----Target
      symbolic-nodename 10.11.1.10 <----- iSCSI host (host 1, not online)

  zone name iscsi-zone-2 vsan 1
    * fcid 0x6d0101 [pwwn 21:00:00:20:37:6f:fe:54] <-----Target
    * fcid 0x6d0201 [pwwn 21:00:00:20:37:a6:a6:5d] <-----Target
      symbolic-nodename iqn.1987-05.com.cisco:01.25589167f74c <-iSCSI host (host 2, not online)
```

Step 13 Bring up the iSCSI hosts (host 1 and host 2).

Step 14 Show all the iSCSI sessions (use the **detail** option for detailed information).

```
switch# show iscsi session
Initiator iqn.1987-05.com.cisco:01.25589167f74c <-----Host 2
Initiator ip addr (s): 10.15.1.11
Session #1
Target iqn.1987-05.com.cisco:05.172.22.92.166.07-05.21000020376ffe54
```



Note The last part of the auto-created target name is the Fibre Channel target's pWWN.

```
VSAN 1, ISID 00023d000001, Status active, no reservation

Session #2
Target iqn.1987-05.com.cisco:05.172.22.92.166.07-05.2100002037a6a65d
VSAN 1, ISID 00023d000001, Status active, no reservation

Initiator 10.11.1.10 <-----Host 1
Initiator name iqn.1987-05.com.cisco:01.e41695d16b1a
Session #1
Target iqn.1987-05.com.cisco:05.172.22.92.166.07-01.21000020376ffd97
VSAN 1, ISID 00023d000001, Status active, no reservation
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Step 15 Verify the details of the two iSCSI initiators.

```
switch# show iscsi initiator
iSCSI Node name is ign.1987-05.com.cisco:01.25589167f74c <-----
  Initiator ip addr (s): 10.15.1.11
  iSCSI alias name: oasis11.cisco.com
  Node WWN is 20:02:00:0b:fd:44:68:c2 (dynamic)
  Member of vsans: 1
  Number of Virtual n_ports: 1
  Virtual Port WWN is 20:03:00:0b:fd:44:68:c2 (dynamic)
    Interface iSCSI 7/5, Portal group tag: 0x304
    VSAN ID 1, FCID 0x6d0300

iSCSI Node name is 10.11.1.10 <-----
  iSCSI Initiator name: ign.1987 - 05.com.cisco:01.e41695d16b1a
  iSCSI alias name: oasis10.cisco.com
  Node WWN is 20:04:00:0b:fd:44:68:c2 (dynamic)
  Member of vsans: 1
  Number of Virtual n_ports: 1
  Virtual Port WWN is 20:05:00:0b:fd:44:68:c2 (dynamic)
    Interface iSCSI 7/1, Portal group tag: 0x300
    VSAN ID 1, FCID 0x6d0301
```

Host 2: Initiator ID based on node name because the initiator is entering iSCSI interface 7/5

Host 1: Initiator ID based on IP address because the initiator is entering iSCSI interface 7/1

Step 16 View the active zone set. The iSCSI initiators' FCIDs are resolved.

```
switch# show zoneset active
zoneset name zoneset-iscsi vsan 1
  zone name iscsi-zone-1 vsan 1
    * fcid 0x6d0001 [pwwn 21:00:00:20:37:6f:fd:97]
    * fcid 0x6d0301 [symbolic-nodename 10.11.1.10] <-----

  zone name iscsi-zone-2 vsan 1
    * fcid 0x6d0101 [pwwn 21:00:00:20:37:6f:fe:54]
    * fcid 0x6d0201 [pwwn 21:00:00:20:37:a6:a6:5d]
    * fcid 0x6d0300 [symbolic-nodename
ign.1987-05.com.cisco:01.25589167f74c] <-----
```

FCID resolved for host 1

FCID for host 2

Step 17 The Fibre Channel name server shows the virtual N ports created for the iSCSI hosts.

```
switch# show fcns database
VSAN 1:
-----
FCID          TYPE  PWWN                                (VENDOR)          FC4-TYPE:FEATURE
-----
0x6d0001      NL    21:00:00:20:37:6f:fd:97 (Seagate)         scsi-fcp:target
0x6d0101      NL    21:00:00:20:37:6f:fe:54 (Seagate)         scsi-fcp:target
0x6d0201      NL    21:00:00:20:37:a6:a6:5d (Seagate)         scsi-fcp:target
0x6d0300      N     20:03:00:0b:fd:44:68:c2 (Cisco)         scsi-fcp:init isc..w
0x6d0301      N     20:05:00:0b:fd:44:68:c2 (Cisco)         scsi-fcp:init isc..w
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Step 18 Verify the detailed output of the iSCSI initiator nodes in the Fibre Channel name server.

```
switch# show fcns database fcid 0x6d0300 detail vsan 1
-----
VSAN:1      FCID:0x6d0300
-----
port-wwn (vendor)      :20:03:00:0b:fd:44:68:c2 (Cisco)
node-wwn                :20:02:00:0b:fd:44:68:c2
class                   :2,3
node-ip-addr            :10.15.1.11    <-----
ipa                     :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw  <-----
symbolic-port-name      :

symbolic-node-name
:iqn.1987-05.com.cisco:01.25589167f74c<-----
port-type                :N
port-ip-addr             :0.0.0.0
fabric-port-wwn          :21:91:00:0b:fd:44:68:c0
hard-addr                :0x000000
Total number of entries = 1
```

**IP address of the iSCSI
host**

iSCSI gateway node

**iSCSI initiator ID is
based on the registered
node name**

```
switch# show fcns database fcid 0x6d0301 detail vsan 1
-----
VSAN:1      FCID:0x6d0301
-----
port-wwn (vendor)      :20:05:00:0b:fd:44:68:c2 (Cisco)
node-wwn                :20:04:00:0b:fd:44:68:c2
class                   :2,3
node-ip-addr            :10.11.1.10
ipa                     :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw  <-----
symbolic-port-name      :

symbolic-node-name      :10.11.1.10  <-----
port-type                :N
port-ip-addr             :0.0.0.0
fabric-port-wwn          :21:81:00:0b:fd:44:68:c0
hard-addr                :0x000000
```

iSCSI gateway node

**iSCSI initiator ID is
based on the IP address
registered in
symbolic-node-name
field**


Send documentation comments to mdsfeedback-doc@cisco.com.

Target Storage Device Requiring LUN Mapping

Sample scenario 2 assumes the following configuration (see [Figure 28-40](#)):

- Access control is based on Fibre Channel zoning.
- There is target-based LUN mapping or LUN masking.
- There is no iSCSI authentication (none).
- The iSCSI initiator is assigned to different VSANs.

Figure 28-40 iSCSI Scenario 2



To configure scenario 2 (see [Figure 28-40](#)), follow these steps:

-
- Step 1** Configure null authentication for all iSCSI hosts.
- ```
switch(config)# iscsi authentication none
```
- Step 2** Configure iSCSI to dynamically import all Fibre Channel targets into the iSCSI SAN using auto-generated iSCSI target names.
- ```
switch(config)# iscsi import target fc
```
- Step 3** Configure the Gigabit Ethernet interface in slot 7 port 1 with an IP address and enable the interface.
- ```
switch(config)# int gigabitethernet 7/1
switch(config-if)# ip address 10.11.1.1 255.255.255.0
switch(config-if)# no shut
```
- Step 4** Configure the iSCSI interface in slot 7 port 1 to identify all dynamic iSCSI initiators by the IP address and enable the interface.
- ```
switch(config)# int iscsi 7/1
switch(config-if)# switchport initiator id ip-address
switch(config-if)# no shut
```

Send documentation comments to mdsfeedback-doc@cisco.com.

- Step 5** Configure the Gigabit Ethernet interface in slot 7 port 5 with the IP address and enable the interface.

```
switch(config)# int gigabitethernet 7/5
switch(config-if)# ip address 10.15.1.1 255.255.255.0
switch(config-if)# no shut
```

- Step 6** Configure the iSCSI interface in slot 7 port 5 to identify all dynamic iSCSI initiators by IP address and enable the interface.

```
switch(config)# int iscsi 7/5
switch(config-if)# switchport initiator id name
switch(config-if)# no shut
```

- Step 7** Add static configuration for each iSCSI initiator.

```
switch(config)# iscsi initiator name iqn.1987-05.com.cisco:01.e41695d16b1a <-----Host 2
switch(config-iscsi-init)# static pwwn system-assign 1
switch(config-iscsi-init)# static nwwn system-assign

switch(config)# iscsi initiator ip address 10.15.1.11 <-----Host 1
switch(config-iscsi-init)# static pwwn system-assigned 1
switch(config-iscsi-init)# vsan 2
```



Note

Host 1 is configured in VSAN 2.

- Step 8** View the configured WWNs.



Note

The WWNs are assigned by the system. The initiators are members of different VSANs.

```
switch# show iscsi initiator configured
iSCSI Node name is iqn.1987-05.com.cisco:01.e41695d16b1a
  Member of vsans: 1
  Node WWN is 20:03:00:0b:fd:44:68:c2
  No. of PWWN: 1
  Port WWN is 20:02:00:0b:fd:44:68:c2

iSCSI Node name is 10.15.1.11
  Member of vsans: 2
  No. of PWWN: 1
  Port WWN is 20:06:00:0b:fd:44:68:c2
```

- Step 9** Create a zone with host 1.

```
switch(config)# zone name iscsi-zone-1 vsan 1
```

- Step 10** Add three members to the zone named *iscsi-zone-1*.



Note

Fibre Channel storage for zone membership for the iSCSI initiator, either the iSCSI symbolic node name or the pWWN, can be used. In this case, the pWWN is persistent.

- The following command is based on the symbolic node name.

```
switch(config-zone)# member symbolic-nodename iqn.1987-05.com.cisco:01.e41695d16b1a
```

- The following command is based on the persistent pWWN assigned to the initiator. You can obtain the pWWN from the **show iscsi initiator** output.

```
switch(config-zone)# member pwwn 20:02:00:0b:fd:44:68:c2
```


Send documentation comments to mdsfeedback-doc@cisco.com.

Step 11 Create a zone with host 2 and two Fibre Channel targets.



Note If the host is in VSAN 2, the Fibre Channel targets and zone must also be in VSAN 2.

Step 12 Activate the zone set in VSAN 2

```
switch(config)# zone name iscsi-zone-2 vsan 2
switch(config)# zoneset activate name iscsi-zoneset-v2 vsan 2
Zoneset activation initiated. check zone status
switch# show zoneset active vsan 2
zoneset name iscsi-zoneset-v2 vsan 2
  zone name iscsi-zone-2 vsan 2
    * fcid 0x750001 [pwwn 21:00:00:20:37:6f:fe:54]
    * fcid 0x750101 [pwwn 21:00:00:20:37:a6:a6:5d]
    pwwn 20:06:00:0b:fd:44:68:c2 <-----Host is not online
```

Step 13 Start the iSCSI clients on both hosts and verify that sessions come up.

Step 14 Display the iSCSI sessions to verify the Fibre Channel target and the configured WWNs.

```
switch# show iscsi session
Initiator iqn.1987-05.com.cisco:01.e41695d16b1a
  Initiator ip addr (s): 10.11.1.10
  Session #1
    Discovery session, ISID 00023d000001, Status active

  Session #2
    Target
    iqn.1987-05.com.cisco:05.172.22.92.166.07-01.21000020376ffd97<---- To Fibre Channel
    VSAN 1, ISID 00023d000001, Status active, no reservation target
```

Step 15 Display the iSCSI initiator to verify the configured nWWN and pWWN.

```
switch# show iscsi initiator
iSCSI Node name is iqn.1987-05.com.cisco:01.e41695d16b1a
  Initiator ip addr (s): 10.11.1.10
  iSCSI alias name: oasis10.cisco.com

  Node WWN is 20:03:00:0b:fd:44:68:c2 (configured)<----- The configured nWWN
  Member of vsans: 1
  Number of Virtual n_ports: 1

  Virtual Port WWN is 20:02:00:0b:fd:44:68:c2 (configured)<---- The configured pWWN
  Interface iSCSI 7/1, Portal group tag: 0x300
  VSAN ID 1, FCID 0x680102
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Step 16 Check the Fibre Channel name server.

```
switch# show fcns database vsan 1
VSAN 1:
-----
FCID      TYPE PWWN                      (VENDOR)  FC4-TYPE:FEATURE
-----
0x680001  NL   21:00:00:20:37:6f:fd:97 (Seagate) scsi-fcp:target

0x680102  N    20:02:00:0b:fd:44:68:c2 (Cisco)  scsi-fcp:init isc..w <---
```

**iSCSI initiator in
name server**

Step 17 Verify the details of the iSCSI initiator's FCID in the name server.

```
switch(config)# show fcns database fcid 0x680102 detail vsan 1
-----
VSAN:1      FCID:0x680102
-----
port-wwn (vendor)      :20:02:00:0b:fd:44:68:c2 (Cisco)
node-wwn               :20:03:00:0b:fd:44:68:c2
class                  :2,3
node-ip-addr           :10.11.1.10
ipa                    :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name     :
symbolic-node-name     :iqn.1987-05.com.cisco:01.e41695d16b1a
port-type              :N
port-ip-addr           :0.0.0.0
fabric-port-wwn        :21:81:00:0b:fd:44:68:c0
iSCSI alias name: oasis10.cisco.com
```

Step 18 Check the Fibre Channel name server

```
switch# show fcns database vsan 1
VSAN 1:
-----
FCID      TYPE PWWN                      (VENDOR)  FC4-TYPE:FEATURE
-----
0x680001  NL   21:00:00:20:37:6f:fd:97 (Seagate) scsi-fcp:target

0x680102  N    20:02:00:0b:fd:44:68:c2 (Cisco)  scsi-fcp:init isc..w <-----
```

**iSCSI
initiator in
name server**

Step 19 Verify the details of the iSCSI initiator's FCID in the name server.

```
switch(config)# show fcns database fcid 0x680102 detail vsan 1
-----
VSAN:1      FCID:0x680102
-----
port-wwn (vendor)      :20:02:00:0b:fd:44:68:c2 (Cisco)
node-wwn               :20:03:00:0b:fd:44:68:c2
class                  :2,3
node-ip-addr           :10.11.1.10
ipa                    :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name     :
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
symbolic-node-name    :iqn.1987-05.com.cisco:01.e41695d16b1a
port-type             :N
port-ip-addr          :0.0.0.0
fabric-port-wwn       :21:81:00:0b:fd:44:68:c0
hard-addr             :0x000000
```

Step 20 Verify that zoning has resolved the FCID for the iSCSI client.

```
switch# show zoneset active vsan 1
zoneset name iscsi-zoneset-v1 vsan 1
  zone name iscsi-zone-1 vsan 1
    * fcid 0x680001 [pwwn 21:00:00:20:37:6f:fd:97]
    * fcid 0x680102 [pwwn 20:02:00:0b:fd:44:68:c2]
```

Step 21 Verify that the second initiator is connected to the two Fibre Channel targets in VSAN 2.

```
switch# show iscsi session initiator 10.15.1.11
Initiator 10.15.1.11
  Initiator name iqn.1987-05.com.cisco:01.25589167f74c
  Session #1
    Target iqn.1987-05.com.cisco:05.172.22.92.166.07-05.21000020376ffe54 <-- Session to
    VSAN 2, ISID 00023d000001, Status active, no reservation                first target

  Session #2
    Target iqn.1987-05.com.cisco:05.172.22.92.166.07-05.2100002037a6a65d <-- Session to
    VSAN 2, ISID 00023d000001, Status active, no reservation                second
                                                                              target

switch# show iscsi initiator
iSCSI Node name is 10.15.1.11 <--- Initiator ID is the IP address
  iSCSI Initiator name: iqn.1987-05.com.cisco:01.25589167f74c
  iSCSI alias name: oasis11.cisco.com

  Node WWN is 20:04:00:0b:fd:44:68:c2 (dynamic) <----- Dynamic
  Member of vsans: 2 <--- vsan membership                    WWN as
  Number of Virtual n_ports: 1                                  static WWN
                                                                              not
                                                                              assigned

  Virtual Port WWN is 20:06:00:0b:fd:44:68:c2 (configured) <----- Static
  Interface iSCSI 7/5, Portal group tag: 0x304                pWWN for
  VSAN ID 2, FCID 0x750200                                     the initiator

switch# show fcns database vsan 2
VSAN 2:
-----
FCID          TYPE  PWWN                                (VENDOR)  FC4-TYPE:FEATURE
-----
0x750001      NL    21:00:00:20:37:6f:fe:54 (Seagate)  scsi-fcp:target
0x750101      NL    21:00:00:20:37:a6:a6:5d (Seagate)  scsi-fcp:target

0x750200      N     20:06:00:0b:fd:44:68:c2 (Cisco)   scsi-fcp:init isc..w <-- iSCSI
Total number of entries = 3                                             initiator
                                                                              entry in
                                                                              name server
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
switch# show fcns database fcid 0x750200 detail vsan 2
```

```
-----
VSAN:2      FCID:0x750200
-----
port-wwn (vendor)      :20:06:00:0b:fd:44:68:c2 (Cisco)
node-wwn               :20:04:00:0b:fd:44:68:c2
class                  :2,3
node-ip-addr           :10.15.1.11
ipa                    :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:scsi-fcp:init iscsi-gw
symbolic-port-name     :
symbolic-node-name     :10.15.1.11
port-type              :N
port-ip-addr           :0.0.0.0
fabric-port-wwn        :21:91:00:0b:fd:44:68:c0
hard-addr              :0x000000
Total number of entries = 1
```

```
switch# show zoneset active vsan 2
```

```
zoneset name iscsi-zoneset-v2 vsan 2
  zone name iscsi-zone-2 vsan 2
    * fcid 0x750001 [pwwn 21:00:00:20:37:6f:fe:54]
    * fcid 0x750101 [pwwn 21:00:00:20:37:a6:a6:5d]

    * fcid 0x750200 [pwwn 20:06:00:0b:fd:44:68:c2] <-----
```

**FCID
resolved for
iSCSI
initiator**

Configuring iSCSI Storage Name Services

iSCSI Storage Name Services (iSNS) allow your existing TCP/IP network to function more effectively as a SAN by automating the discovery, management, and configuration of iSCSI devices. To facilitate these functions, the iSNS server and client function as follows:

- The iSNS client registers iSCSI portals and all iSCSI devices accessible through them with an iSNS server.
- The iSNS server provides the following services for iSNS client:
 - Device registration
 - State change notification
 - Remote domain discovery services

All iSCSI devices (both initiator and target), acting as iSNS clients, can register with an iSNS server. iSCSI initiators can then query the iSNS server for a list of targets. The iSNS server will respond with a list of targets that the querying client can access based on configured access control parameters.

In Cisco MDS SAN-OS Release 1.3, a Cisco MDS 9000 Family switch can act as an iSNS client and register all available iSCSI targets with an external iSNS server. Cisco MDS SAN-OS Release 2.0(1b), and later, supports iSNS server functionality on all switches in the Cisco MDS 9000 Family with IPS modules or MPS-14/2 modules installed. This allows external iSNS clients, such as an iSCSI initiator, to register with the switch and discover all available iSCSI targets in the SAN.

Send documentation comments to mdsfeedback-doc@cisco.com.

iSNS Client Functionality

The iSNS client functionality on each IPS interface (Gigabit Ethernet interface or subinterface or PortChannel) registers information with an iSNS server. You must specify an iSNS server's IP address by creating an iSNS profile, adding the server's IP address to it, then assign (or "tag") the profile to the interface. An iSNS profile can be tagged to one or more interfaces.

Once a profile is tagged to an interface, the switch opens a TCP connection to the iSNS server IP address (using the well-known iSNS port number 3205) in the profile and registers network entity and portal objects; a unique entity is associated with each IPS interface. The switch then searches the Fibre Channel name server (FCNS) database and switch configuration to find storage nodes to register with the iSNS server.

Statically mapped virtual targets are registered if the associated Fibre Channel pWWN is present in the FCNS database and no access control configuration prevents it. A dynamically mapped target is registered if dynamic target importing is enabled. See the [“Presenting Fibre Channel Targets as iSCSI Targets” section on page 28-54](#) for more details on how iSCSI imports Fibre Channel targets.

A storage node is deregistered from the iSNS server when it becomes unavailable when configuration changes (such as access control change or dynamic import disabling) or the Fibre Channel storage port goes off-line. It is registered again when the node comes back online.

When the iSNS client is unable to register or deregister objects with the iSNS server (for example, the client is unable to make a TCP connection to the iSNS server), it retries every minute to reregister all iSNS objects for the affected interfaces with the iSNS server. The iSNS client uses a registration interval value of 15 minutes. If the client fails to refresh the registration during this interval, the server will deregister the entries.

Untagging a profile also causes the network entity and portal to be deregistered from that interface.

Creating an iSNS Profile

To create an iSNS profile, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# isns profile name MyIsns switch(config-isns-profile)# | Creates a profile called MyIsns. |
| Step 3 | switch(config-isns-profile)# server 10.10.100.211 | Specifies an iSNS server IP address for this profile. |
| | switch(config-isns-profile)# no server 10.20.100.211 | Removes a configured iSNS server from this profile. |

To remove an iSNS profile, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# no isns profile name OldIsns | Removes a configured iSNS profile called OldIsns. |

Send documentation comments to mdsfeedback-doc@cisco.com.

To tag a profile to an interface, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# conf t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# interface gigabitethernet 4/1 switch(config-if)# | Configures the specified Gigabit Ethernet interface. |
| Step 3 | switch(config-if)# isns MyIsns | Tags a profile to an interface. |

To untag a profile from an interface, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# conf t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# interface gigabitethernet 5/1 switch(config-if)# | Configures the specified Gigabit Ethernet interface. |
| Step 3 | switch(config-if)# no isns OldIsns | Untags a profile from an interface. |

Use the **isns reregister** command in EXEC mode to re-register associated iSNS objects with the iSNS server.

```
switch# isns reregister gigabitethernet 1/4
switch# isns reregister port-channel 1
```

Verifying iSNS Client Configuration

Use the **show isns profile** command to view configured iSNS profiles. Profile ABC has two portals registered with the iSNS server. Each portal corresponds to a particular interface. Profile XYZ has a specified iSNS server, but does not have any tagged interfaces configured (see [Example 28-45](#) and [Example 28-46](#)).

Example 28-45 Displays Information for Configured iSNS Profiles

```
switch# show isns profile
iSNS profile name ABC
tagged interface GigabitEthernet2/3
tagged interface GigabitEthernet2/2
iSNS Server 10.10.100.204

iSNS profile name XYZ
iSNS Server 10.10.100.211
```

Example 28-46 Displays a Specified iSNS Profile

```
switch# show isns profile ABC
iSNS profile name ABC
tagged interface GigabitEthernet2/3
tagged interface GigabitEthernet2/2
iSNS Server 10.10.100.204
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Use the **show isns profile counters** command to view all configured profiles with the iSNS PDU statistics for each tagged interface (see [Example 28-47](#) and [Example 28-48](#)).

Example 28-47 Displays Configured Profiles with iSNS Statistics

```
switch# show isns profile counters
iSNS profile name ABC
tagged interface port-channel 1
iSNS statistics
  Input 54 pdus (registration/deregistration pdus only)
    Reg pdus 37, Dereg pdus 17
  Output 54 pdus (registration/deregistration pdus only)
    Reg pdus 37, Dereg pdus 17
iSNS Server 10.10.100.204

iSNS profile name XYZ
tagged interface port-channel 2
iSNS statistics
  Input 30 pdus (registration/deregistration pdus only)
    Reg pdus 29, Dereg pdus 1
  Output 30 pdus (registration/deregistration pdus only)
    Reg pdus 29, Dereg pdus 1
iSNS Server 10.1.4.218
```

Example 28-48 Displays iSNS Statistics for a Specified Profile

```
switch# show isns profile ABC counters
iSNS profile name ABC
tagged interface port-channel 1
iSNS statistics
  Input 54 pdus (registration/deregistration pdus only)
    Reg pdus 37, Dereg pdus 17
  Output 54 pdus (registration/deregistration pdus only)
    Reg pdus 37, Dereg pdus 17
iSNS Server 10.10.100.204
```

Use the **show isns** command to view all objects registered on the iSNS server and specified in the given profile (see [Example 28-49](#)).

Example 28-49 Displays iSNS Queries

```
switch# show isns query ABC gigabitethernet 2/3
iSNS server: 10.10.100.204
Init: iqn.1991-05.com.w2k
  Alias: <MS SW iSCSI Initiator>
Tgt : iqn.1987-05.com.cisco:05.172.22.94.22.02-03
Tgt : iqn.1987-05.com.cisco:05.172.22.94.22.02-03.210000203762fa34
  nWWN: 200000203762fa34
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Use the **show interface** command to view the iSNS profile to which an interface is tagged (see [Example 28-50](#)).

Example 28-50 Displays Tagged iSNS Interfaces

```
switch# show interface gigabitethernet 2/3
GigabitEthernet2/3 is up
Hardware is GigabitEthernet, address is 0005.3000.ae94
Internet address is 10.10.100.201/24
MTU 1500 bytes
Port mode is IPS
Speed is 1 Gbps
Beacon is turned off
Auto-Negotiation is turned on
iSNS profile ABC
^^^^^^^^^^^^^^^^
5 minutes input rate 112 bits/sec, 14 bytes/sec, 0 frames/sec
5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
1935 packets input, 132567 bytes
  4 multicast frames, 0 compressed
  0 input errors, 0 frame, 0 overrun 0 fifo
1 packets output, 42 bytes, 0 underruns
  0 output errors, 0 collisions, 0 fifo
  0 carrier errors
```

iSNS Server Functionality

When enabled, the iSNS server on the Cisco 9000 Family MDS switch tracks all registered iSCSI devices. As a result, iSNS clients can locate other iSNS clients by querying the iSNS server. The iSNS server also provides the following functionalities:


- Allows iSNS clients to register, deregister, and query other iSNS clients registered with the iSNS server.
- Provides centralized management for enforcing access control to provide or deny access to targets from specific initiators.
- Provides a notification mechanism for registered iSNS clients to receive change notifications on the status change of other iSNS clients.
- Provides a single access control configuration for both Fibre Channel and iSCSI devices.
- Discovers iSCSI targets that do not have direct IP connectivity to the iSCSI initiators.

Example Scenario

The iSNS server provides uniform access control across Fibre Channel and iSCSI devices by utilizing both Fibre Channel zoning information and iSCSI access control information and configuration. An iSCSI initiator acting as an iSNS client only discovers devices it is allowed to access based on both sets of access control information. [Figure 28-41](#) provides an example of this scenario.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 28-41 Using iSNS Servers in the Cisco MDS Environment



In [Figure 28-41](#), iqn.host1 and iqn.host2 are iSCSI initiators. P1 and P2 are Fibre Channel targets. The two initiators are in different zones: Zone 1 consists of iqn.host1 and target P1, and Zone 2 consists of iqn.host2 and target P2. iSNS server functionality is enabled on both switches, SW-1 and SW-2. The registration process proceeds as follows:

1. Initiator iqn.host1 registers with SW-1, port GigabitEthernet2/1.
2. Initiator iqn.host2 registers with SW-2, port GigabitEthernet3/1.
3. Initiator iqn.host1 issues an iSNS query to SW-1 to determine all accessible targets.
4. The iSNS server in turn queries the Fibre Channel Name Server (FCNS) to obtain a list of devices that are accessible (that is, in the same zone) by the query originator. This query yields only P1.
5. The iSNS server then queries its own database to convert the Fibre Channel devices to the corresponding iSCSI targets. This is based on the iSCSI configuration, such as virtual-target and its access control setting or whether the dynamic Fibre Channel target import feature is enabled or disabled.
6. The iSNS server sends a response back to the query initiator. This response contains a list all iSCSI portals known to the iSNS server. This means iqn.host1 can choose to login to target P1 via either SW-1 (at GigabitEthernet 2/1) or SW-2 (at GigabitEthernet 3/1).
7. If the initiator chooses to login to SW-1 and later that port becomes inaccessible (for example, GigabitEthernet 2/1 goes down), the initiator has the choice to move to connect to target P1 via port GigabitEthernet 3/1 on SW-2 instead.
8. If the target either goes down or is removed from the zone, the iSNS server sends out an iSNS State Change Notification (SCN) message to the initiator so that the initiator can remove the session.

Send documentation comments to mdsfeedback-doc@cisco.com.

Enabling the iSNS Server

Before iSNS server feature can be enabled, iSCSI must be enabled (see the [“Enabling iSCSI” section on page 28-53](#)). When you disable iSCSI, iSNS is automatically disabled. When the iSNS server is enabled on a switch, every IPS port whose corresponding iSCSI interface is up is capable of servicing iSNS registration and query requests from external iSNS clients.

To enable the iSNS server, follow these steps:

| | Command | Purpose |
|--------|--|-------------------------------------|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# isns-server enable | Enables the iSNS server. |
| | switch(config)# no isns-server enable | Disables (default) the iSNS server. |

iSCSI Configuration Distribution

You can use the CFS infrastructure to distribute the iSCSI initiator configuration to iSNS servers across fabric. This allows iSNS server running on any switch to provide to a querying iSNS client a list of iSCSI devices available anywhere on the fabric. For information on CFS, see the [Chapter 9, “Using the CFS Infrastructure.”](#)

To enable configuration distribution, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# isns distribute | Uses the CFS infrastructure to distribute the iSCSI virtual target configuration to all switches in the fabric. |
| | switch(config)# no isns distribute | Stops (default) the distribution of iSCSI virtual target configuration to all switches in the fabric. |

ESI Retry Count Configuration

The iSNS server periodically queries for Entity Status Inquiry (ESI) port from the iSNS clients. Receipt of a response indicates that the client is still alive. If the client fails to respond, after a configured number of retries (the default is 3), the client is deregistered from the server.

The ESI interval value cannot be configured; The minimum allowed value is 60 seconds; iSNS clients may indicate a higher value at registration. An ESI interval value of 0 is permitted and is used by the iSCSI client to indicate to the server that it does not want to be monitored using ESI. In such cases, the registrations made by the client will remain valid till explicitly deregistered or if the iSNS server feature is disabled.

To configure the ESI retry count, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# isns esi retries 6 | Configures the ESI to retry contacting the client up to 6 times. |

Send documentation comments to mdsfeedback-doc@cisco.com.

iSNS Client Registration and Deregistration

An iSNS client cannot query the iSNS server until it has registered. You can use the **show isns database** command to display all registered iSNS clients and their associated configuration.

iSNS client deregistration can occur either explicitly or when the iSNS server detects that it can no longer reach the client (through ESI monitoring).

iSNS client registration and deregistration result in status change notifications (SCNs) being generated to all interested iSNS clients.

Target Discovery

iSCSI initiators discover targets by issuing queries to the iSNS server. The server supports *DevGetNext* requests to search the list of targets and *DevAttrQuery* to determine target and portal details, such as the IP address or port number to which to connect.

On receiving a query request from the iSCSI client, the iSNS server queries the Fibre Channel Name Server (FCNS) to obtain a list of Fibre Channel targets that are accessible by the querying initiator. The result of this query depends on zoning configuration currently active and current configuration(s) of the initiator. The iSNS server will subsequently use the iSCSI target configuration(s) (virtual target and dynamic import configuration) to translate the Fibre Channel target to an equivalent iSCSI target. At this stage it also applies any access control configured for the virtual target. A response message with the target details is then sent back to the query initiator.

The iSNS server sends a consolidated response containing all possible targets and portals to the querying initiator. For example, if a Fibre Channel target is exported as different iSCSI targets on different IPS interfaces, the iSNS server will respond with a list of all possible iSCSI targets and portals.

In order to keep the list of targets updated, the iSNS server sends state change notifications (SCN) to the client whenever an iSCSI target becomes reachable or unreachable. The client is then expected to rediscover its list of accessible targets by initiating another iSNS query. Reachability of iSCSI targets changes when any one of the following occurs:

1. Target goes up or down
2. Dynamic import of FC target configuration changes
3. Zone set changes
4. Default zone access control changes
5. IPS interface state changes
6. Initiator configuration that change make the target accessible or inaccessible.

Send documentation comments to mdsfeedback-doc@cisco.com.

Verifying the iSNS Server Configuration

Use the **show isns config** command to view the ESI interval and the summary information about the iSNS database contents (see [Example 28-51](#)).

Example 28-51 Displays the iSNS Server Configuration of ESI Interval and Database Contents

```
switch# show isns config
Server Name: switch1(Cisco Systems) Up since: Fri Jul 30 04:08:16 2004
  Index: 1      Version: 1      TCP Port: 3205
  fabric distribute (remote sync): ON
  ESI
    Non Response Threshold: 5 Interval(seconds): 60
  Database contents
    Number of Entities: 2
    Number of Portals: 3
    Number of iSCSI devices: 4
    Number of Portal Groups: 0
```

Use the **show isns database** command to view detailed information about the contents of the iSNS database (see [Example 28-52](#) through [Example 28-55](#)). This command displays the full iSNS database giving all the entities, nodes, and portals registered in the database. This command without options only displays explicitly registered objects. The asterisk next to the VSAN ID indicates that the iSCSI node is in the default zone for that VSAN.

Example 28-52 Displays Explicitly Registered Objects

```
switch# show isns database
Entity Id: dp-204
  Index: 2      Last accessed: Fri Jul 30 04:08:46 2004

iSCSI Node Name: ign.1991-05.comdp-2041
  Entity Index: 2
  Node Type: Initiator(2)      Node Index: 0x1
  SCN Bitmap: OBJ_UPDATED|OBJ_ADDED|OBJ_REMOVED|TARGET&SELF
  Node Alias: <MS SW iSCSI Initiator>

  VSANS: 1(*), 5(*)
Portal IP Address: 192.168.100.2      TCP Port: 4179
  Entity Index: 2      Portal Index: 1
  ESI Interval: 0      ESI Port: 4180      SCN Port: 4180
```

[Example 28-53](#) displays information about both virtual and registered iSCSI initiators/targets.

Example 28-53 Displays the Full Database With Both Registered and Configured Nodes and Portals

```
switch# show isns database full
Entity Id: isns.entity.mds9000
  Index: 1      Last accessed: Fri Jul 30 04:08:16 2004

iSCSI Node Name: ign.com.cisco.disk1
  Entity Index: 1
  Node Type: Target(1)      Node Index: 0x80000001
  WWN(s):
    22:00:00:20:37:39:dc:45

  VSANS:
iSCSI Node Name: ign.isns-first-virtual-target
  Entity Index: 1
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```

Node Type: Target(1)          Node Index: 0x80000002

VSANS:
iSCSI Node Name: ign.com.cisco.disk2
Entity Index: 1
Node Type: Target(1)          Node Index: 0x80000003
WWN(s):
    22:00:00:20:37:39:dc:45

VSANS:
Portal IP Address: 192.168.100.5      TCP Port: 3205
Entity Index: 1      Portal Index: 3

Portal IP Address: 192.168.100.6      TCP Port: 3205
Entity Index: 1      Portal Index: 5

Entity Id: dp-204
Index: 2          Last accessed: Fri Jul 30 04:08:46 2004

iSCSI Node Name: ign.1991-05.com.microsoft:dp-2041
Entity Index: 2
Node Type: Initiator(2)          Node Index: 0x1
SCN Bitmap: OBJ_UPDATED|OBJ ADDED|OBJ REMOVED|TARGET&SELF
Node Alias: <MS SW iSCSI Initiator>

VSANS: 1(*), 5(*)
Portal IP Address: 192.168.100.2      TCP Port: 4179
Entity Index: 2      Portal Index: 1
ESI Interval: 0      ESI Port: 4180      SCN Port: 4180

```

Example 28-54 displays the virtual targets entries on the current switch.



Note

The **local** option is only available for virtual targets.

Example 28-54 Displays the Virtual Target Information in the Local Switch

```

switch# show isns database virtual-targets local
Entity Id: isns.entity.mds9000
Index: 1          Last accessed: Fri Jul 30 04:08:16 2004

iSCSI Node Name: ign.com.cisco.disk1
Entity Index: 1
Node Type: Target(1)          Node Index: 0x80000001
WWN(s):
    22:00:00:20:37:39:dc:45

VSANS:
iSCSI Node Name: ign.isns-first-virtual-target
Entity Index: 1
Node Type: Target(1)          Node Index: 0x80000002

VSANS:
iSCSI Node Name: ign.com.cisco.disk2
Entity Index: 1
Node Type: Target(1)          Node Index: 0x80000003
WWN(s):
    22:00:00:20:37:39:dc:45

VSANS:
Portal IP Address: 192.168.100.5      TCP Port: 3205
Entity Index: 1      Portal Index: 3

```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
Portal IP Address: 192.168.100.6      TCP Port: 3205
Entity Index: 1      Portal Index: 5
```

[Example 28-55](#) provides the virtual target information for a specific remote switch. The remote switch is specified using the switch ID (the WWN of the switch).

Example 28-55 Displays Virtual Target for a Specified Switch

```
switch# show isns database virtual-targets switch 20:00:00:0d:ec:01:04:40
Entity Id: isns.entity.mds9000
  Index: 1      Last accessed: Fri Jul 30 04:08:16 2004

iSCSI Node Name: ign.com.cisco.disk1
  Entity Index: 1
  Node Type: Target(1)      Node Index: 0x80000001
  WWN(s):
    22:00:00:20:37:39:dc:45

  VSANS:
iSCSI Node Name: ign.isns-first-virtual-target
  Entity Index: 1
  Node Type: Target(1)      Node Index: 0x80000002

  VSANS:
iSCSI Node Name: ign.com.cisco.disk2
  Entity Index: 1
  Node Type: Target(1)      Node Index: 0x80000003
  WWN(s):
    22:00:00:20:37:39:dc:45

  VSANS:
Portal IP Address: 192.168.100.5      TCP Port: 3205
Entity Index: 1      Portal Index: 3

Portal IP Address: 192.168.100.6      TCP Port: 3205
Entity Index: 1      Portal Index: 5
```

Use the **show isns node** command to display attributes of nodes registered with the iSNS server (see [Example 28-56](#) through [Example 28-58](#)). If you do not specify any options the server displays the name and node type attribute in a compact format; one per line.

Example 28-56 Displays Explicitly Registered Objects

```
switch# show isns node all
-----
iSCSI Node Name                                     Type
-----
ign.1987-05.com.cisco:05.switch1.02-03.22000020375a6c8      Target
...
ign.com.cisco.disk1                                         Target
ign.com.cisco.ipdisk                                       Target
ign.isns-first-virtual-target                             Target
ign.1991-05.cw22                                           Target
ign.1991-05.cw53                                           Target
```

Example 28-57 Displays of the Specified Node

```
switch# show isns node name ign.com.cisco.disk1
iSCSI Node Name: ign.com.cisco.disk1
Entity Index: 1
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
Node Type: Target(1)      Node Index: 0x80000001
WWN(s):
    22:00:00:20:37:39:dc:45
VSANS: 1
```

Example 28-58 Displays the Attribute Details for All Nodes

```
switch# show isns node all detail
iSCSI Node Name: ign.1987-05.com.cisco:05.switch1.02-03.22000020375a6c8f
Entity Index: 1
Node Type: Target(1)      Node Index: 0x30000003
Configured Switch WWN: 20:00:00:0d:ec:01:04:40
WWN(s):
    22:00:00:20:37:5a:6c:8f
VSANS: 1
...
iSCSI Node Name: ign.com.cisco.disk1
Entity Index: 1
Node Type: Target(1)      Node Index: 0x80000001
Configured Switch WWN: 20:00:00:0d:ec:01:04:40
WWN(s):
    22:00:00:20:37:39:dc:45
VSANS: 1

iSCSI Node Name: ign.com.cisco.ipdisk
Entity Index: 1
Node Type: Target(1)      Node Index: 0x80000002
Configured Switch WWN: 20:00:00:0d:ec:01:04:40
WWN(s):
    22:00:00:20:37:5a:70:1a
VSANS: 1

iSCSI Node Name: ign.isns-first-virtual-target
Entity Index: 1
Node Type: Target(1)      Node Index: 0x80000003
Configured Switch WWN: 20:00:00:0d:ec:01:04:40

iSCSI Node Name: ign.parna.121212
Entity Index: 1
Node Type: Target(1)      Node Index: 0x80000004
Configured Switch WWN: 20:00:00:0d:ec:01:04:40

iSCSI Node Name: ign.parna.121213
Entity Index: 1
Node Type: Target(1)      Node Index: 0x80000005
Configured Switch WWN: 20:00:00:0d:ec:01:04:40
```

Use the **show isns portal** command to display the attributes of a portal along with its accessible nodes (see [Example 28-59](#) through [Example 28-63](#)). You can specify portals by using the switch WWN-interface combination or the IP address-port number combination.

Example 28-59 Displays the Attribute Information for All Portals

```
switch# show isns portal all
```

| IPAddress | TCP Port | Index | SCN Port | ESI port |
|---------------|----------|-------|----------|----------|
| 192.168.100.5 | 3205 | 3 | - | - |
| 192.168.100.6 | 3205 | 5 | - | - |

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 28-60 Displays Detailed Attribute Information for All Portals

```
switch# show isns portal all detail
Portal IP Address: 192.168.100.5      TCP Port: 3205
      Entity Index: 1      Portal Index: 3

Portal IP Address: 192.168.100.6      TCP Port: 3205
      Entity Index: 1      Portal Index: 5
```

Example 28-61 Displays Virtual Portals

```
switch# show isns portal virtual
```

| IPAddress | TCP Port | Index | SCN Port | ESI | port |
|---------------|----------|-------|----------|-----|------|
| 192.168.100.5 | 3205 | 3 | - | - | |
| 192.168.100.6 | 3205 | 5 | - | - | |

Example 28-62 Displays Virtual Portals for the Specified Switch

```
switch# show isns portal virtual switch 20:00:00:0d:ec:01:04:40
```

| IPAddress | TCP Port | Index | SCN Port | ESI | port |
|---------------|----------|-------|----------|-----|------|
| 192.168.100.5 | 3205 | 3 | - | - | |
| 192.168.100.6 | 3205 | 5 | - | - | |

Example 28-63 Displays Detailed Information for the Virtual Portals in the Specified Switch

```
switch# show isns portal virtual switch 20:00:00:0d:ec:01:04:40 detail
Portal IP Address: 192.168.100.5      TCP Port: 3205
      Entity Index: 1      Portal Index: 3
      Switch WWN: 20:00:00:0d:ec:01:04:40
      Interface: GigabitEthernet2/3

Portal IP Address: 192.168.100.6      TCP Port: 3205
      Entity Index: 1      Portal Index: 5
      Switch WWN: 20:00:00:0d:ec:01:04:40
      Interface: GigabitEthernet2/5
```

Use the **show isns entity** command to display the attributes of an entity along with the list of portals and nodes in that entity (see [Example 28-64](#) through [Example 28-68](#)). If you do not specify any option, this command displays the entity ID and number of nodes or portals associated with the entity in a compact format; one per line.

Example 28-64 Displays All Registered Entries

```
switch1# show isns entity
```

| Entity ID | Last Accessed |
|-----------|-------------------------|
| dp-204 | Tue Sep 7 23:15:42 2004 |

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 28-65 Displays All Entities in the Database

```
switch# show isns entity all
```

```
-----
Entity ID                                     Last Accessed
-----
isns.entity.mds9000                          Tue Sep  7 21:33:23 2004
dp-204                                       Tue Sep  7 23:15:42 2004
```

Example 28-66 Displays the Entity with the Specified ID

```
switch1# show isns entity id dp-204
Entity Id: dp-204
      Index: 2                Last accessed: Tue Sep  7 23:15:42 2004
```

Example 28-67 Displays Detailed Information for All Entities in the Database

```
switch1# show isns entity all detail
Entity Id: isns.entity.mds9000
      Index: 1                Last accessed: Tue Sep  7 21:33:23 2004

Entity Id: dp-204
      Index: 2                Last accessed: Tue Sep  7 23:16:34 2004
```

Example 28-68 Displays Virtual Entities

```
switch# show isns entity virtual
Entity Id: isns.entity.mds9000
      Index: 1                Last accessed: Thu Aug  5 00:58:50 2004

Entity Id: dp-204
      Index: 2                Last accessed: Thu Aug  5 01:00:23 2004
```

Use the **show iscsi global config** command to display information about import targets (see [Example 28-69](#) and [Example 28-70](#)).

Example 28-69 Displays the Import Target Settings for the Specified Switch

```
switch# show isns iscsi global config switch 20:00:00:05:ec:01:04:00
iSCSI Global configuration:
      Switch: 20:00:00:05:ec:01:04:00 iSCSI Auto Import: Enabled
```

Example 28-70 Displays the Import Target Settings for All Switches

```
switch# show isns iscsi global config all
iSCSI Global configuration:
      Switch: 20:00:44:0d:ec:01:02:40 iSCSI Auto Import: Enabled
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Use the **show cfs peers** command to display CFS peers switch information about the iSNS application (see [Example 28-71](#)).

Example 28-71 Displays the CFS Peer Switch Information for the iSNS Application

```
switch# show cfs peers name isns

Scope      : Physical
-----
Switch WWN          IP Address
-----
20:00:00:00:ec:01:00:40  10.10.100.11  [Local]

Total number of entries = 1
```

IPS Module Core Dumps

IPS core dumps are different from the system's kernel core dumps for other modules. When the IPS module's operating system (OS) unexpectedly resets, it is useful to obtain a copy of the memory image (called a IPS core dump) to identify the cause of the reset. Under that condition, the IPS module sends the core dump to the supervisor module for storage. Cisco MDS switches have two levels of IPS core dumps:

- Partial core dumps (default)—Each partial core dump consists of four parts (four files). All four files are saved in the active supervisor module.

Use the **show cores** command to list these files.

- Full core dumps—Each full core dump consists of 75 parts (75 files). The IPS core dumps for the MPS-14/2 module and the Cisco MDS 9216i Switch only contains 38 parts. This dump cannot be saved on the supervisor module due to its large space requirement. They are copied directly to an external TFTP server.

Use the **system cores tftp:** command to configure an external TFTP server to copy the IPS core dump (and other core dumps).

To configure IPS core dumps on the IPS module, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ips core dump full ips core dump full' successfully set for module 9 | Configures a dump of the full core generation for all IPS modules in the switch. |
| | switch(config)# no ips core dump full ips core dump partial' successfully set for module 9 | Configures a dump of the partial core generation for the IPS module in slot 9. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Default Settings

Table 28-3 lists the default settings for Gigabit Ethernet parameters.

Table 28-3 Default Gigabit Ethernet Parameters

| Parameters | Default |
|-------------------|-----------------------------------|
| IP MTU frame size | 1500 bytes for all Ethernet ports |
| Auto-negotiation | Enabled. |
| Promiscuous mode | Disabled |

Table 28-4 lists the default settings for FCIP parameters.

Table 28-4 Default FCIP Parameters

| Parameters | Default |
|---|-----------------------------|
| TCP default port for FCIP | 3225. |
| minimum-retransmit-time | 200 ms. |
| keepalive-timeout | 60 seconds. |
| max-retransmissions | 4 retransmissions. |
| PMTU discovery | Enabled. |
| pmtu-enable reset-timeout | 3600 seconds. |
| SACK | Enabled. |
| max-bandwidth | 1Gbps. |
| min-available-bandwidth | 500 Mbps. |
| round-trip-time | 1 ms. |
| buffer size | 0 KB. |
| Control TCP and data connection | No packets are transmitted. |
| TCP congestion window monitoring | Enabled. |
| Burst size | 50KB. |
| TCP connection mode | Active mode is enabled. |
| special-frame | Disabled. |
| FCIP timestamp | Disabled. |
| acceptable-diff range to accept packets | + or - 2000 ms. |
| B port keepalive responses | Disabled. |
| Write acceleration | Disabled. |
| Tape acceleration | Disabled. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 28-5 lists the default settings for iSCSI parameters.

Table 28-5 Default iSCSI Parameters

| Parameters | Default |
|--|--|
| Number of TCP connections | One per iSCSI session. |
| minimum-retransmit-time | 300 ms. |
| keepalive-timeout | 60 seconds. |
| max-retransmissions | 4 retransmissions. |
| PMTU discovery | Enabled. |
| pmtu-enable reset-timeout | 3600 seconds. |
| SACK | Enabled. |
| max-bandwidth | 1 G.bps |
| min-available-bandwidth | 70 Mbps. |
| round-trip-time | 1 ms. |
| buffer size | 4096 KB. |
| Control TCP and data connection | No packets are transmitted. |
| TCP congestion window monitoring | Enabled. |
| Burst size | 50KB. |
| Jitter | 500 Microseconds. |
| TCP connection mode | Active mode is enabled. |
| Fibre Channel targets to iSCSI | Not imported. |
| Advertising iSCSI target | Advertised on all Gigabit Ethernet interfaces, subinterfaces, PortChannel interfaces, and PortChannel subinterfaces |
| iSCSI hosts mapping to virtual Fibre Channel hosts | Dynamic mapping. |
| Dynamic iSCSI initiators | Members of the VSAN 1. |
| Identifying initiators | iSCSI node names. |
| Advertising static virtual targets | No initiators are allowed to access a virtual target (unless explicitly configured). |
| iSCSI login authentication | CHAP or none authentication mechanism. |
| revert-primary-port | Disabled. |
| Header and data digest | Enabled automatically when iSCSI initiators send requests. This feature cannot be configured and is not available in store-and-forward mode. |
| iSNS registration interval | 60 seconds (not configurable). |
| iSNS registration interval retries | 3. |
| Fabric distribution | Enabled. |



Configuring IPsec Network Security

IP Security (IPsec) Protocol is a framework of open standards that provides data confidentiality, data integrity, and data authentication between participating peers. It is developed by the Internet Engineering Task Force (IETF). IPsec provides security services at the IP layer, including protecting one or more data flows between a pair of hosts, between a pair of security gateways, or between a security gateway and a host. The overall IPsec implementation is per the latest version of RFC2401. Cisco SAN-OS IPsec implements RFC 2402 through RFC 2410.

Refer to the following website for further information on the IPsec RFCs:
<http://www.ietf.org/html.charters/ipsec-charter.html>.

IPsec uses the Internet Key Exchange (IKE) protocol to handle protocol and algorithm negotiation and to generate the encryption and authentication keys to be used by IPsec. While IKE can be used with other protocols, its initial implementation is with the IPsec protocol. IKE provides authentication of the IPsec peers, negotiates IPsec security associations, and establishes IPsec keys. IKE uses RFCs 2408, 2409, 2410, 2412, and additionally, implements the draft-ietf-ipsec-ikev2-16.txt draft.

Refer to the following website for further information on the IKE draft:
<http://www.ietf.org/html.charters/ipsec-charter.html>



Note

The term IPsec is sometimes used to describe the entire protocol of IPsec data services and IKE security protocols and is also sometimes used to describe only the data services.

This chapter includes the following sections:

- [About IPsec, page 29-2](#)
- [About IKE, page 29-3](#)
- [IPsec Prerequisites, page 29-3](#)
- [IPsec Compatibility, page 29-4](#)
- [IPsec and IKE Terminology, page 29-5](#)
- [Supported IPsec Transforms and Algorithms, page 29-6](#)
- [Supported IKE Transforms and Algorithms, page 29-6](#)
- [Initializing IKE, page 29-7](#)
- [Configuring the IKE Domain, page 29-7](#)
- [About IKE Tunnels, page 29-8](#)
- [IKE Policy Negotiation, page 29-8](#)
- [Clearing IKE Tunnels or Domains, page 29-11](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

- [Refreshing SAs, page 29-12](#)
- [Configuring IPsec, page 29-12](#)
- [IPsec Maintenance, page 29-23](#)
- [Global Lifetime Values, page 29-23](#)
- [Displaying IKE Configurations, page 29-24](#)
- [Displaying IPsec Configurations, page 29-25](#)
- [Sample FCIP Configuration, page 29-30](#)
- [Sample iSCSI Configuration, page 29-34](#)
- [Default Settings, page 29-36](#)

About IPsec

IPsec provides security for transmission of sensitive information over unprotected networks such as the Internet. IPsec acts at the network layer, protecting and authenticating IP packets between participating IPsec devices (peers).

IPsec provides the following network security services. In general, the local security policy dictates the use of one or more of these services between two participating IPsec devices:

- Data confidentiality—The IPsec sender can encrypt packets before transmitting them across a network.
- Data integrity—The IPsec receiver can authenticate packets sent by the IPsec sender to ensure that the data has not been altered during transmission.
- Data origin authentication—The IPsec receiver can authenticate the source of the IPsec packets sent. This service is dependent upon the data integrity service.
- Anti-replay protection—The IPsec receiver can detect and reject replayed packets.



Note

The term data authentication is generally used to mean data integrity and data origin authentication. Within this chapter it also includes anti-replay services, unless otherwise specified.

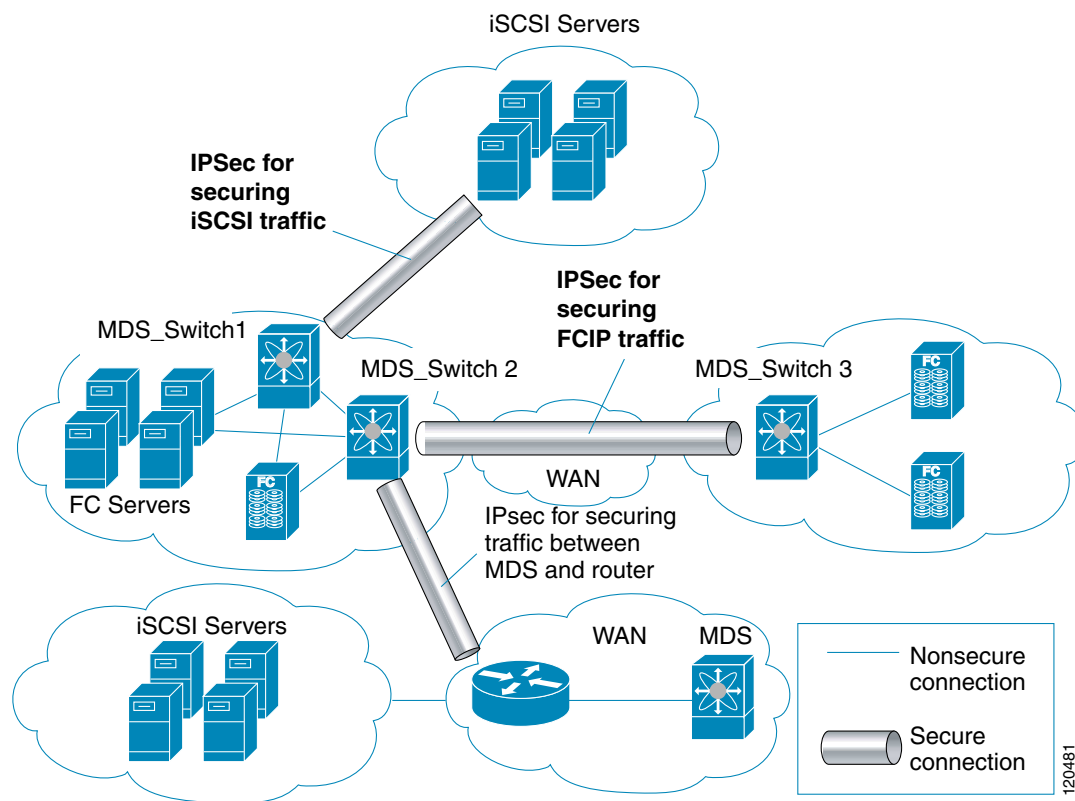
With IPsec, data can be transmitted across a public network without fear of observation, modification, or spoofing. This enables applications such as Virtual Private Networks (VPNs), including intranets, extranets, and remote user access.

IPsec as implemented in Cisco SAN-OS software supports the Encapsulating Security Payload (ESP) protocol. This protocol encapsulates the data to be protected and provides data privacy services, optional data authentication, and optional anti-replay services.

[Figure 29-1](#) shows different IPsec scenarios.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 29-1 FCIP and iSCSI Scenarios Using MPS-14-2 Modules



About IKE

IKE automatically negotiates IPsec security associations and generates keys for all switches using the IPsec feature. Specifically, IKE provides these benefits:

- Allows you to refresh IPsec SAs.
- Allows IPsec to provide anti-replay services.
- Supports a manageable, scalable IPsec configuration.
- Allows dynamic authentication of peers.

IPsec Prerequisites

To use the IPsec feature, you need to perform the following tasks:

- Obtain the ENTERPRISE_PKG license and/or the SAN Extension over IP license. The Enterprise package enables IPsec for iSCSI and the SAN Extension over IP package enables IPsec for FCIP.(see [Chapter 3, “Obtaining and Installing Licenses”](#)).
- Configure IKE as described in the [“Initializing IKE” section on page 29-7](#).

Send documentation comments to mdsfeedback-doc@cisco.com.

**Note**

The IPsec feature inserts new headers in existing packets. (see [“Configuring the MTU Frame Size” section on page 28-6](#) for more information).

IPsec Compatibility

IPsec features are compatible with the following Cisco MDS hardware running Cisco MDS SAN-OS Release 2.0(1b) or later:

- MPS-14/2 modules in Cisco MDS 9200 Switches or Cisco MDS 9500 Directors
- Cisco MDS 9216i Switch with the 14/2-Port multiprotocol capability in the integrated supervisor module. Refer to the *Cisco MDS 9200 Series Hardware Installation Guide* for more information on the Cisco MDS 9216i Switch.
- The IPsec feature is not supported on the management interface.

IPsec features are compatible with the following fabric set up:

- Two connected Cisco MDS 9200 Switches or Cisco MDS 9500 Directors running Cisco MDS SAN-OS Release 2.0(1b) or later.
- A Cisco MDS 9200 Switches or Cisco MDS 9500 Directors running Cisco MDS SAN-OS Release 2.0(1b) or later connected to any IPsec compliant device.
- The following features are not supported in the Cisco SAN-OS implementation of the IPsec feature:
 - Authentication Header (AH).
 - Transport mode.
 - Security association bundling.
 - Manually configuring security associations.
 - Per host security association option in a crypto map.
 - Security association idle timeout
 - Dynamic crypto maps.

**Note**

Any reference to crypto maps in this document, only refers to static crypto maps.

Send documentation comments to mdsfeedback-doc@cisco.com.

IPsec and IKE Terminology

The terms used in this chapter are explained in this section.

- Security association (SA)—An agreement between two participating peers on the entries required to encrypt and decrypt IP packets. Two SAs are required for each peer in each direction (inbound and outbound) to establish bidirectional communication between the peers. Sets of bidirectional SA records are stored in the SA database (SAD). IPsec uses IKE to negotiate and bring up SAs. Each SA record includes the following information:
 - Security parameter index (SPI)—A number which, together with a destination IP address and security protocol, uniquely identifies a particular SA. When using IKE to establish the SAs, the SPI for each SA is a pseudo-randomly derived number.
 - Peer—A switch or other device that participates in IPsec. For example, a Cisco MDS switch or other Cisco routers that support IPsec.
 - Transform—A list of operations done to provide data authentication and data confidentiality. For example, one transform is the ESP protocol with the HMAC-MD5 authentication algorithm.
 - Session key—The key used by the transform to provide security services.
 - Lifetime—A lifetime counter (in seconds and bytes) is maintained from the time the SA is created. When the time limit expires the SA is no longer operational and, if required, is automatically renegotiated (rekeyed).
 - Mode of operation—Two modes of operation are generally available for IPsec: tunnel mode and transport mode. The Cisco SAN-OS implementation of IPsec only supports the tunnel mode. The IPsec tunnel mode encrypts and authenticates the IP packet, including its header. The gateways encrypt traffic on behalf of the hosts and subnets. The Cisco SAN-OS implementation of IPsec does not support transport mode.



Note

The term *tunnel mode* is different from the term *tunnel* used to indicate secure communication path between two peers, such as two switches connected by an FCIP link.

- Anti-replay—A security service where the receiver can reject old or duplicate packets in order to protect itself against replay attacks. IPsec provides this optional service by use of a sequence number combined with the use of data authentication.
- Data authentication—Data authentication can refer either to integrity alone or to both integrity and authentication (data origin authentication is dependent on data integrity).
 - Data integrity—Verifies that data has not been altered.
 - Data origin authentication—Verifies that the data was actually sent by the claimed sender.
- Data confidentiality—A security service where the protected data cannot be observed.
- Data flow—A grouping of traffic, identified by a combination of source address/mask, destination address/mask, IP next protocol field, and source and destination ports, where the protocol and port fields can have the values of any. Traffic matching a specific combination of these values is logically grouped together into a data flow. A data flow can represent a single TCP connection between two hosts, or it can represent traffic between two subnets. IPsec protection is applied to data flows.

Send documentation comments to mdsfeedback-doc@cisco.com.

- Perfect forward secrecy (PFS)—A cryptographic characteristic associated with a derived shared secret value. With PFS, if one key is compromised, previous and subsequent keys are not compromised, because subsequent keys are not derived from previous keys.
- Security Policy Database (SPD)—an ordered list of policies applied to traffic. A policy decides if a packet requires IPsec processing, if should be allowed in clear text, or if it should be dropped.
 - The IPsec SPDs are derived from user configuration of crypto maps.
 - The IKE SPD is configured by the user.

Supported IPsec Transforms and Algorithms

The component technologies implemented for IPsec include the following transforms:

- Advanced Encrypted Standard (AES) is an encryption algorithm. It implements either 128 or 256 bits using Cipher Block Chaining (CBC) or counter mode.
- Data Encryption Standard (DES) is used to encrypt packet data and implements the mandatory 56-bit DES-CBC. CBC requires an initialization vector (IV) to start encryption. The IV is explicitly given in the IPsec packet.
- Triple DES (3DES) is a stronger form of DES with 168-bit encryption keys that allow sensitive information to be transmitted over untrusted networks.



Note

Cisco SAN-OS images with strong encryption are subject to United States government export controls, and have a limited distribution. Images to be installed outside the United States require an export license. Customer orders might be denied or subject to delay due to United States government regulations. Contact your sales representative or distributor for more information, or send e-mail to export@cisco.com.

- Message Digest 5 (MD5) is a hash algorithm with the HMAC variant. HMAC is a keyed hash variant used to authenticate data.
- Secure Hash Algorithm (SHA-1) is a hash algorithm with the Hash Message Authentication Code (HMAC) variant.
- AES-XCBC-MAC is a Message Authentication Code (MAC) using the AES algorithm.

Supported IKE Transforms and Algorithms

The component technologies implemented for IKE include the following transforms:

- Diffie-Hellman (DH) is a public-key cryptography protocol which allows two parties to establish a shared secret over an unsecure communications channel. Diffie-Hellman is used within IKE to establish session keys. Group 1 (768-bit), Group 2 (1024-bit), and Group 5 (1536-bit) groups are supported.
- Advanced Encrypted Standard (AES) is an encryption algorithm. It implements either 128 bits using Cipher Block Chaining (CBC) or counter mode.
- Data Encryption Standard (DES) is used to encrypt packet data and implements the mandatory 56-bit DES-CBC. CBC requires an initialization vector (IV) to start encryption. The IV is explicitly given in the IPsec packet.

Send documentation comments to mdsfeedback-doc@cisco.com.

- Triple DES (3DES) is a stronger form of DES with 168-bit encryption keys that allow sensitive information to be transmitted over untrusted networks.



Note

Cisco SAN-OS images with strong encryption are subject to United States government export controls, and have a limited distribution. Images to be installed outside the United States require an export license. Customer orders might be denied or subject to delay due to United States government regulations. Contact your sales representative or distributor for more information, or send e-mail to export@cisco.com.

- Message Digest 5 (MD5) is a hash algorithm with the HMAC variant. HMAC is a keyed hash variant used to authenticate data.
- Secure Hash Algorithm (SHA-1) is a hash algorithm with the Hash Message Authentication Code (HMAC) variant.
- The switch authentication algorithm uses the preshared keys based on the IP address (see [“Setting the Global Preshared Key”](#) section on page 19-6 for more information on preshared keys).

Initializing IKE

The IKE feature must first be enabled and configured so the IPsec feature can establish data flow with the required peer.

You cannot disable IKE if IPsec is enabled. When you disable the IKE feature, the IKE configuration is cleared from the running configuration.

To enable IKE, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# crypto ike enable switch(config)# no crypto ike enable | Enables the IKE feature. Disables (default) the IKE feature. |

Configuring the IKE Domain

You must apply the IKE configurations to an IPsec domain to allow traffic to reach the supervisor module in the local switch.

To configure the IPsec domain, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# crypto ike domain ipsec switch(config-ike-ipsec)# | Allows IKE configurations for IPsec domains. |

[Send documentation comments to mdsfeedback-doc@cisco.com.](mailto:mdsfeedback-doc@cisco.com)

About IKE Tunnels

An IKE tunnel is a secure IKE session between two end points. IKE creates this tunnel to protect IKE messages used in IPsec SA negotiations.

You can display the IKE tunnels using the **show crypto ike domain ipsec sa** command in EXEC mode.

Two versions of IKE are used in the Cisco SAN-OS implementation.

- IKE version 1 (IKEv1) is implemented using RFC 2407, 2408, 2409, and 2412.
- IKE version 2 (IKEv2) is a simplified and more efficient version and does not interoperate with IKEv1. IKEv2 is implemented using the draft-ietf-ipsec-ikev2-16.txt draft.

IKE Policy Negotiation

To protect IKE negotiations, each IKE negotiation begins with a common (shared) IKE policy. An IKE policy defines a combination of security parameters to be used during the IKE negotiation. By default, no IKE policy is configured. You must create IKE policies at each peer. This policy states which security parameters will be used to protect subsequent IKE negotiations and mandates how peers are authenticated. You can create multiple, prioritized policies at each peer to ensure that at least one policy will match a remote peer's policy.

You can configure the policy based on the encryption algorithm (DES, 3DES, or AES), the hash algorithm (SHA or MD5), and the DH group (1, 2, or 5). Each policy can contain a different combination of parameter values. A unique priority number identifies the configured policy. This number ranges from 1 (highest priority) to 255 (lowest priority). You can create multiple policies in a switch. If you need to connect to a remote peer, you must ascertain that at least one policy in the local switch contains the identical parameter values configured in the remote peer. If several policies have identical parameter configurations, the policy with the lowest number is selected.

[Table 29-1](#) provides a list of allowed transform combinations.

Table 29-1 *IKE Transform Configuration Parameters*

| Parameter | Accepted Values | Keyword | Default Value |
|-----------------------|----------------------|------------------|----------------|
| encryption algorithm | 56-bit DES-CBC | des | 3des |
| | 168-bit DES | 3des | |
| | 128-bit AES | aes | |
| hash algorithm | SHA-1 (HMAC variant) | sha | sha |
| | MD5 (HMAC variant) | md5 | |
| authentication method | Preshared keys | Not configurable | Preshared keys |
| DH group identifier | 768-bit DH | 1 | 1 |
| | 1024-bit DH | 2 | |
| | 1536-bit DH | 5 | |

Send documentation comments to mdsfeedback-doc@cisco.com.

**Note**

The following table lists the supported and verified settings for IPsec and IKE encryption authentication algorithms on the Microsoft Windows and Linux platforms:

| Platform | IKE | IPsec |
|---|-----------------------------------|-------------|
| Microsoft iSCSI initiator, Microsoft IPsec implementation on Microsoft Windows 2000 platform | 3DES, SHA-1 or MD5, DH group 2 | 3DES, SHA-1 |
| Cisco iSCSI initiator, Free Swan IPsec implementation on Linux platform | 3DES, MD5, DH group 1 | 3DES, MD5 |

**Note**

When you configure the hash algorithm, the corresponding HMAC version is used as the authentication algorithm.

When the IKE negotiation begins, IKE looks for an IKE policy that is the same on both peers. The peer that initiates the negotiation will send all its policies to the remote peer, and the remote peer will try to find a match. The remote peer looks for a match by comparing its own highest priority policy against the other peer's received policies. The remote peer checks each of its policies in order of its priority (highest priority first) until a match is found.

A match is found when the two peers have the same encryption, hash algorithm, authentication algorithm, and DH group values. If a match is found, IKE completes the security negotiation and the IPsec SAs are created.

If an acceptable match is not found, IKE refuses negotiation and the IPsec data flows will not be established.

To configure the IKE negotiation parameters, follow these steps:

| | Command | Purpose |
|---------------|---|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# crypto ike domain ipsec switch(config-ike-ipsec)# | Allows IPsec domains to be configured in this switch. |
| Step 3 | switch(config-ike-ipsec)# key Sample address 10.10.100.232 | Sets the preshared key for the specified peer. |
| | switch(config-ike-ipsec)# no key Sample address 10.10.100.232 | Deletes the preshared key for the specified peer. |
| Step 4 | switch(config-ike-ipsec)# policy 1 switch(config-ike-ipsec-policy)# | Identifies the policy to be configured. |
| | switch(config-ike-ipsec)# no policy 1 | Deletes the identified policy. |
| Step 5 | switch(config-ike-ipsec-policy)# encryption des | Configures the encryption policy. |
| | switch(config-ike-ipsec-policy)# no encryption aes | Defaults to 3DES encryption. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|---|--------------------------------|
| Step 6 | switch(config-ike-ipsec-policy)# group 5 | Configures the DH group. |
| | switch(config-ike-ipsec-policy)# no group 2 | Defaults to DH group 1. |
| Step 7 | switch(config-ike-ipsec-policy)# hash md5 | Configures the hash algorithm. |
| | switch(config-ike-ipsec-policy)# no hash md5 | Defaults to SHA. |

Optional Configurations

You can optionally configure the following parameters for the IKE feature:

- The lifetime association within each policy—The lifetime ranges from 600 to 86,400 seconds. The default is 86,400 seconds (equals one day).
- The keepalive time for each peer if you use IKEv2—The keepalive ranges from 120 to 86,400 seconds. The default is 3,600 seconds (equals one hour).
- The initiator version for each peer—IKE v1 or IKE v2 (default). Your choice of initiator version does not affect interoperability when the remote device initiates the negotiation. Configure this option if the peer device supports IKEv1 and you can play the initiator role for IKE with the specified device.



Caution

You may need to configure this option even when the switch doesn't behave as an IKE initiator under normal circumstances. Always using this option guarantees a faster recovery of traffic flows in case of failures.



Tip

The keepalive time only applies to IKEV2 peers and not to all peers.



Note

When IPSec implementations in the host prefer to initiate the IPSec rekey, be sure to configure the IPSec lifetime value in the Cisco MDS switch to be higher than the lifetime value in the host.

To configure the lifetime association for each policy, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# crypto ike domain ipsec switch(config-ike-ipsec)# | Allows IPsec domains to be configured in this switch. |
| Step 3 | switch(config-ike-ipsec)# policy 1 switch(config-ike-ipsec-policy)# | Identifies the policy to be configured. |
| Step 4 | switch(config-ike-ipsec-policy) lifetime seconds 6000 | Configures a lifetime of 6,000 seconds. |
| | switch(config-ike-ipsec-policy)# no lifetime seconds 6000 | Deletes the configured lifetime value and defaults to 86,400 seconds. |

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure the keepalive time for each peer, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# crypto ike domain ipsec switch(config-ike-ipsec)# | Allows IPsec domains to be configured in this switch. |
| Step 3 | switch(config-ike-ipsec)# keepalive 60000 | Configures the keepalive time for all peers to be 60,000 seconds. |
| | switch(config-ike-ipsec)# no keepalive 60000 | Deletes the configured keepalive time and defaults to 3,600 seconds. |

To configure the initiator version, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# crypto ike domain ipsec switch(config-ike-ipsec)# | Allows IPsec domains to be configured in this switch. |
| Step 3 | switch(config-ike-ipsec)# initiator version 1 address 10.10.10.1 | Configures the switch to use IKEv1 when initiating IKE with Device 10.10.10.0 |
| | switch(config-ike-ipsec)# no initiator version 1 address 10.10.10.0 | Defaults to IKEv2 for the specified device. |
| | switch(config-ike-ipsec)# no initiator version 1 | Defaults to IKEv2 for all devices. |

Clearing IKE Tunnels or Domains

If a IKE tunnel ID is not specified for the IKE configuration, you can clear all existing IKE domain connections by issuing the **clear crypto ike domain ipsec sa** command in EXEC mode.

```
switch# clear crypto ike domain ipsec sa
```



Caution

When you delete all the SAs within a specific IKEv2 tunnel, then that IKE tunnel is automatically deleted.

If an SA is specified for the IKE configuration, you can clear the specified IKE tunnel ID connection by issuing the **clear crypto ike domain ipsec sa IKE_tunnel-ID** command in EXEC mode.

```
switch# clear crypto ike domain ipsec sa 51
```



Caution

When you delete the IKEv2 tunnel, the associated IPsec tunnel under that IKE tunnel is automatically deleted.

Send documentation comments to mdsfeedback-doc@cisco.com.

Refreshing SAs

Use the **crypto ike domain ipsec rekey sa *sa-index*** command to refresh the SAs after performing IKEv2 configuration changes.

Configuring IPsec

IPsec provides secure data flows between participating peers. Multiple IPsec data flows can exist between two peers to secure different data flows, with each tunnel using a separate set of SAs.

After you have completed IKE configuration, configure IPsec.

To configure IPsec in each participating IPsec peer, follow these steps:

-
- | | |
|---------------|--|
| Step 1 | Identify the peers for the traffic to which secure tunnels should be established. |
| Step 2 | Configure the transform set with the required protocols and algorithms. |
| Step 3 | Create the crypto map and apply Access Control Lists (ACLs), transform set, peer, lifetime values as applicable. |
| Step 4 | Apply the crypto map to the required interface. |
-

Crypto ACLs

IP Access Control Lists (IP-ACLs) provide basic network security to all switches in the Cisco MDS 9000 Family. IP-ACLs restrict IP-related traffic based on the configured IP filters. Refer to the [“IP Access Control Lists” section on page 26-5](#) for details on creating and defining IP-ACLs.

In the context of crypto maps, ACLs are different from regular ACLs. Regular ACLs determine what traffic to forward or block at an interface. For example, ACLs can be created to protect all IP traffic between Subnet A and Subnet Y or Telnet traffic between Host A and Host B.

Crypto ACLs are used to define which IP traffic requires crypto protection and which traffic does not.

Crypto ACLs associated with IPsec crypto map entries have four primary functions:

- Select outbound traffic to be protected by IPsec (permit = protect).
- Indicate the data flow to be protected by the new SAs (specified by a single permit entry) when initiating negotiations for IPsec SAs.
- Process inbound traffic to filter out and discard traffic that should have been protected by IPsec.
- Determine whether or not to accept requests for IPsec SAs on behalf of the requested data flows when processing IKE negotiation from the IPsec peer.



Tip

If you want some traffic to receive one type of IPsec protection (for example, encryption only) and other traffic to receive a different type of IPsec protection (for example, both authentication and encryption), create two ACLs. Use both ACLs in different crypto maps to specify different IPsec policies.

Send documentation comments to mdsfeedback-doc@cisco.com.

To create ACLs, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# ip access-list List1 permit ip 10.1.1.100 0.0.0.255 11.1.1.100 0.0.0.255 | Permits all IP traffic from and to the specified networks. |



Note

The **show access-list** command does not display the crypto map entries. Use the **show crypto map** command to display the associated entries.

Add permit and deny statements as appropriate (see the “[IP Access Control Lists](#)” section on page 26-5). Each permit and deny specifies conditions to determine which IP packets must be protected.

Crypto ACL Guidelines

Follow these guidelines when configuring ACLs for the IPsec feature:

- The Cisco SAN-OS software only allows name-based IP ACLs.
- When an IP ACL is applied to a crypto map, the following applies:
 - Permit—applying the IPsec feature to the traffic.
 - Deny—allowing clear text (default).



Note

IKE traffic (UDP port 500) is implicitly transmitted in clear text.

- The IPsec feature only considers the source and destination IP addresses and subnet masks.



Note

The IPsec feature ignores the port numbers and protocol fields.

- The **permit** option causes all IP traffic that matches the specified conditions to be protected by crypto, using the policy described by the corresponding crypto map entry.
- The **deny** option prevents traffic from being protected by crypto. The first deny statement causes the traffic to be in clear text.
- The crypto ACL you define is applied to an interface after you define the corresponding crypto map entry and apply the crypto map set to the interface.
- Different ACLs must be used in different entries of the same crypto map set.
- Inbound and outbound traffic is evaluated against the same outbound IPsec ACL. Therefore, the ACL's criteria is applied in the forward direction to traffic exiting your switch, and the reverse direction to traffic entering your switch.

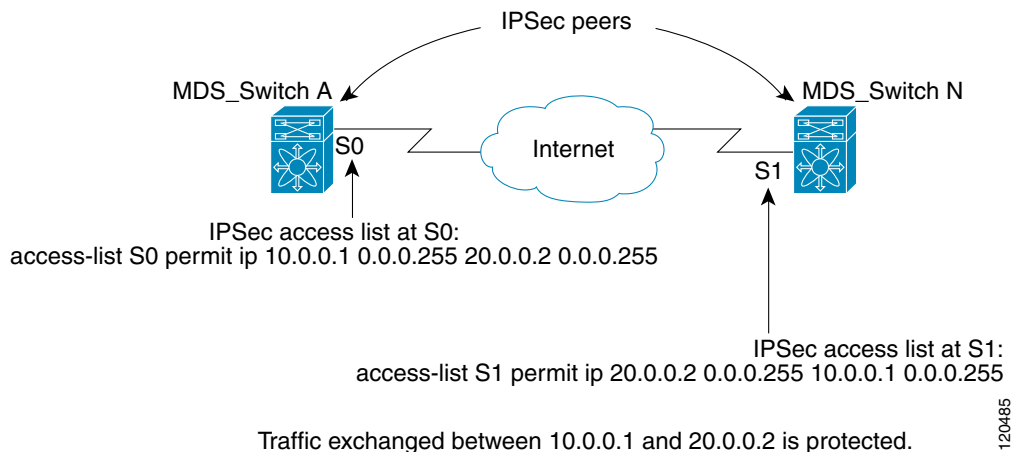
Send documentation comments to mdsfeedback-doc@cisco.com.

- Each ACL filter assigned to the crypto map entry is equivalent to one security policy entry. The IPsec feature supports up to 120 security policy entries for each MPS-14/2 module and Cisco MDS 9216i Switch.
- In Figure 29-2, IPsec protection is applied to traffic between switch interface S0 (IP address 10.0.0.1) and switch interface S1 (IP address 20.0.0.2) as the data exits switch A's S0 interface enroute to switch interface S1. For traffic from 10.0.0.1 to 20.0.0.2, the ACL entry on switch A is evaluated as follows:
 - source = IP address 10.0.0.1
 - dest = IP address 20.0.0.2

For traffic from 20.0.0.2 to 10.0.0.1, that same ACL entry on switch A is evaluated as follows:

- source = IP address 20.0.0.2
- dest = IP address 10.0.0.1

Figure 29-2 IPsec Processing of Crypto ACLS



- If you configure multiple statements for a given crypto ACL which is used for IPsec, the first permit statement that is matched is used to determine the scope of the IPsec SA. Later, if traffic matches a different permit statement of the crypto ACL, a new, separate IPsec SA is negotiated to protect traffic matching the newly matched ACL statement.
- Unprotected inbound traffic that matches a permit entry in the crypto ACL for a crypto map entry flagged as IPsec is dropped, because this traffic was expected to be protected by IPsec.
- Use the **show ip access-lists** command to view all IP ACLs. The IP ACLs used for traffic filtering purposes are also used for crypto.

Mirror Image Crypto ACLs

For every crypto ACL specified for a crypto map entry defined at the local peer, define a mirror image crypto ACL at the remote peer. This configuration ensures that IPsec traffic applied locally can be processed correctly at the remote peer.



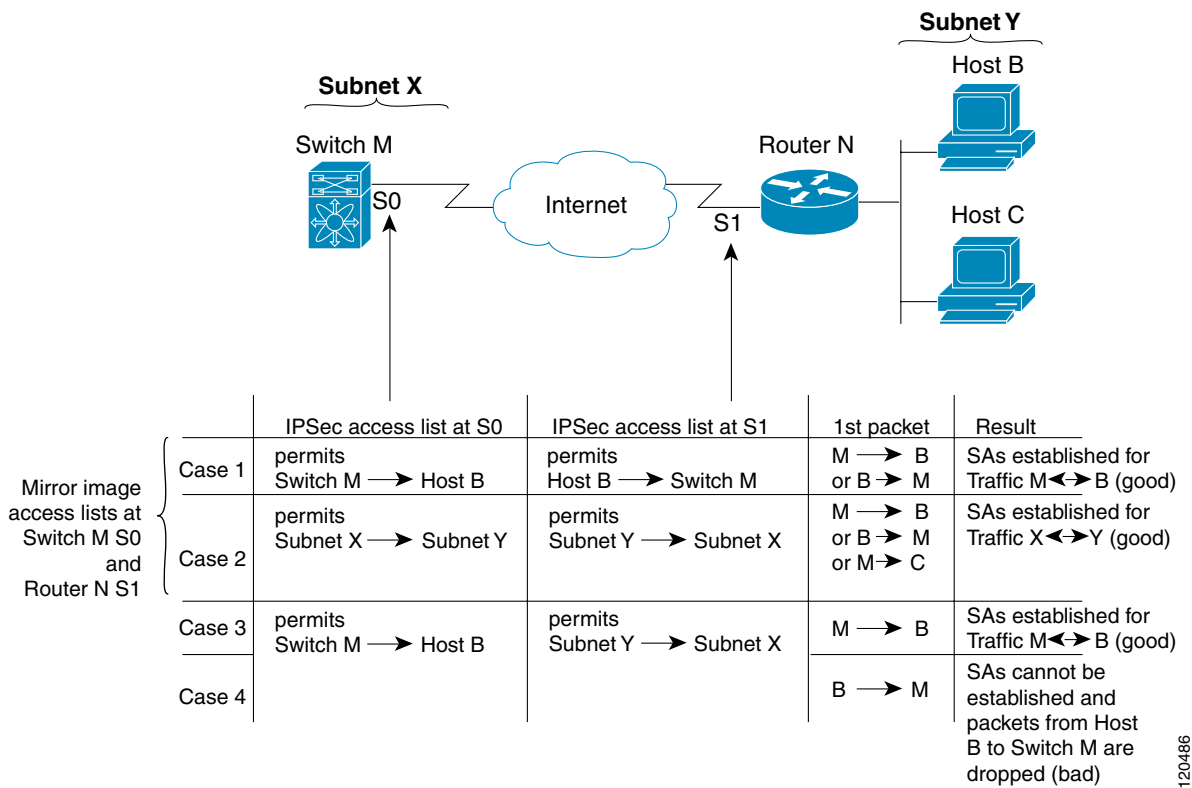
Tip

The crypto map entries themselves must also support common transforms and must refer to the other system as a peer.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 29-3 shows some sample scenarios with and without mirror image ACLs.

Figure 29-3 IPsec Processing of Mirror Image Configuration



As Figure 29-3 indicates, IPsec SAs (SAs) can be established as expected whenever the two peers' crypto ACLs are mirror images of each other. However, an IPsec SA can be established only some of the time when the ACLs are not mirror images of each other. This can happen in the case where an entry in one peer's ACL is a subset of an entry in the other peer's ACL, such as shown in Cases 3 and 4 of Figure 29-3. IPsec SA establishment is critical to IPsec—without SAs, IPsec does not work, causing any packets matching the crypto ACL criteria to be silently dropped instead of being forwarded with IPsec security.

In Figure 29-3, an SA cannot be established in Case 4. This is because SAs are always requested according to the crypto ACLs at the initiating packet's end. In Case 4, router N requests that all traffic between Subnet X and Subnet Y be protected, but this is a superset of the specific flows permitted by the crypto ACL at switch M so the request is therefore not permitted. Case 3 works because switch M's request is a subset of the specific flows permitted by the crypto ACL at router N.

Because of the complexities introduced when crypto ACLs are not configured as mirror images at peer IPsec devices, Cisco strongly encourages you to use mirror image crypto ACLs.

Send documentation comments to mdsfeedback-doc@cisco.com.

The any Keyword in Crypto ACLs



Tip

We recommend that you configure mirror image crypto ACLs for use by IPsec and that you avoid using the **any** option.

The **any** option in a permit statement is discouraged when you have multicast traffic flowing through the IPsec interface—this configuration can cause multicast traffic to fail.

The **permit any any** statement causes all outbound traffic to be protected (and all protected traffic sent to the peer specified in the corresponding crypto map entry) and requires protection for all inbound traffic. Then, all inbound packets that lack IPsec protection are silently dropped, including packets for routing protocols, NTP, echo, echo response, and so forth.

You need to be sure you define which packets to protect. If you must use the **any** option in a permit statement, you must preface that statement with a series of deny statements to filter out any traffic (that would otherwise fall within that permit statement) that you do not want to be protected.

Transform Sets in IPsec

A transform set represents a certain combination of security protocols and algorithms. During the IPsec security association negotiation, the peers agree to use a particular transform set for protecting a particular data flow.

You can specify multiple transform sets, and then specify one or more of these transform sets in a crypto map entry. The transform set defined in the crypto map entry is used in the IPsec security association negotiation to protect the data flows specified by that crypto map entry's access list.

During IPsec security association negotiations with IKE, the peers search for a transform set that is the same at both peers. When such a transform set is found, it is selected and applied to the protected traffic as part of both peers' IPsec security associations.



Tip

If you change a transform set definition, the change is only applied to crypto map entries that reference the transform set. The change is not applied to existing security associations, but used in subsequent negotiations to establish new security associations. If you want the new settings to take effect sooner, you can clear all or part of the security association database.



Note

When you enable IPsec, the Cisco SAN-OS software automatically creates a default transform set (ipsec_default_tranform_set) using AES-128 encryption and SHA-1 authentication algorithms.

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 29-2 provides a list of allowed transform combinations for IPsec.

Table 29-2 IPsec Transform Configuration Parameters

| Parameter | Accepted Values | Keyword |
|--|------------------------------|-------------------------|
| encryption algorithm | 56-bit DES-CBC | esp-des |
| | 168-bit DES | esp-3des |
| | 128-bit AES-CBC | esp-aes 128 |
| | 128-bit AES-CTR ¹ | esp-aes 128 ctr |
| | 256-bit AES-CBC | esp-aes 256 |
| | 256-bit AES-CTR ¹ | esp-aes 256 ctr |
| hash/authentication algorithm ¹ (optional) | SHA-1 (HMAC variant) | esp-sha1-hmac |
| | MD5 (HMAC variant) | esp-md5-hmac |
| | AES-XCBC-MAC | esp-aes-xcbc-mac |

1. If you configure the AES counter (CTR) mode, you must also configure the authentication algorithm.



Note

The following table lists the supported and verified settings for IPsec and IKE encryption authentication algorithms on the Microsoft Windows and Linux platforms:

| Platform | IKE | IPsec |
|---|-----------------------------------|-------------|
| Microsoft iSCSI initiator, Microsoft IPsec implementation on Microsoft Windows 2000 platform | 3DES, SHA-1 or MD5, DH group 2 | 3DES, SHA-1 |
| Cisco iSCSI initiator, Free Swan IPsec implementation on Linux platform | 3DES, MD5, DH group 1 | 3DES, MD5 |

To configure transform sets, follow these steps:

| | Command | Purpose |
|---------------|--|---|
| Step 1 | <code>switch# config terminal</code> <code>switch(config)#</code> | Enters configuration mode. |
| Step 2 | <code>switch(config)# crypto transform-set</code> <code>domain ipsec test esp-3des esp-md5-hmac</code> | Configures a transform set called test specifying the 3DES encryption algorithm and the MD5 authentication algorithm. Refer to Table 29-2 to verify the allowed transform combinations. |
| | <code>switch(config)# no crypto transform-set</code> <code>domain ipsec test esp-3des esp-md5-hmac</code> | Deletes the applied transform set. |
| | <code>switch(config)# crypto transform-set</code> <code>domain ipsec test esp-3des</code> | Configures a transform set called test specifying the 3DES encryption algorithm. In this case, the default no authentication is performed. |
| | <code>switch(config)# no crypto transform-set</code> <code>domain ipsec test esp-3des</code> | Deletes the applied transform set. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Crypto Map Entries

Once you have created the crypto ACLs and transform sets, you can create crypto map entries that combine the various parts of the IPsec SA, including the following:

- The traffic to be protected by IPsec (per the crypto ACL). A crypto map set can contain multiple entries, each with a different ACL.
- The granularity of the flow to be protected by a set of SAs.
- The IPsec-protected traffic destination (who the remote IPsec peer is).
- The local address to be used for the IPsec traffic (applying to an interface).
- The IPsec security to be applied to this traffic (selecting from a list of one or more transform sets).
- Other parameters to define an IPsec SA.

Crypto map entries with the same crypto map name (but different map sequence numbers) are grouped into a crypto map set.

When you apply a crypto map set to an interface, the following events occur:

- A Security Policy Database (SPD) is created for that interface.
- All IP traffic passing through the interface is evaluated against the SPD.

If a crypto map entry sees outbound IP traffic that requires protection, an SA is negotiated with the remote peer according to the parameters included in the crypto map entry.

The policy derived from the crypto map entries is used during the negotiation of SAs. If the local switch initiates the negotiation, it will use the policy specified in the crypto map entries to create the offer to be sent to the specified IPsec peer. If the IPsec peer initiates the negotiation, the local switch checks the policy from the crypto map entries and decide whether to accept or reject the peer's request (offer).

For IPsec to succeed between two IPsec peers, both peers' crypto map entries must contain compatible configuration statements.

SA Establishment Between Peers

When two peers try to establish an SA, they must each have at least one crypto map entry that is compatible with one of the other peer's crypto map entries.

For two crypto map entries to be compatible, they must at least meet the following criteria:

- The crypto map entries must contain compatible crypto ACLs (for example, mirror image ACLs). If the responding peer entry is in the local crypto, the ACL must be permitted by the peer's crypto ACL.
- The crypto map entries must each identify the other peer or must have auto peer configured.
- If you create more than one crypto map entry for a given interface, use the `seq-num` of each map entry to rank the map entries: the lower the `seq-num`, the higher the priority. At the interface that has the crypto map set, traffic is evaluated against higher priority map entries first.
- The crypto map entries must have at least one transform set in common where IKE negotiations are carried out and SAs are established. During the IPsec SA negotiation, the peers agree to use a particular transform set when protecting a particular data flow.

When a packet matches a permit entry in a particular ACL, the corresponding crypto map entry is tagged, and connections are established.

Send documentation comments to mdsfeedback-doc@cisco.com.

Crypto Map Configuration Guidelines

When configuring crypto map entries, follow these guidelines:

- The sequence number for each crypto map decides the order in which the policies are applied. A lower sequence number is assigned a higher priority.
- Only one ACL is allowed for each crypto map entry (the ACL itself can have multiple permit or deny entries).
- When the tunnel endpoint is the same as the destination address, you can use the **auto-peer** option to dynamically configure the peer.

To create mandatory crypto map entries, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| | switch(config)# crypto map domain ipsec SampleMap 31 ips-hacl1(config-crypto-map-ip)# | Place you in the crypto map configuration mode for the entry named SampleMap with 31 as its sequence number. |
| Step 2 | switch(config)# no crypto map domain ipsec SampleMap 3 | Deletes the specified crypto map entry. |
| | switch(config)# no crypto map domain ipsec SampleMap | Deletes the entire crypto map set called SampleMap. |
| | switch(config-crypto-map-ip)# match address SampleAcl | Names a ACL to determine which traffic should be protected and not protected by IPsec in the context of this crypto map entry. |
| | switch(config-crypto-map-ip)# no match address SampleAcl | Deletes the matched address. |
| Step 4 | switch(config-crypto-map-ip)# set peer 10.1.1.1 | Configures a specific peer IP address. |
| | switch(config-crypto-map-ip)# no set peer 10.1.1.1 | Deletes the configured peer. |
| Step 5 | switch(config-crypto-map-ip)# set transform-set SampleTransform1 SampleTransformfor2 | Specifies which transform sets are allowed for the specified crypto map entry or entries. List multiple transform sets in order of priority (highest priority first). |
| | switch(config-(crypto-map-ip))# no set transform-set | Deletes the association of all the transform sets (regardless of you specifying a transform set name). |

SA Lifetime Negotiation

You can override the global lifetime values (size and time) by configuring a SA-specific lifetime value.

To specify SA lifetime negotiation values, you can optionally configure the lifetime value for a specified crypto map. If you do, this value overrides the globally set values. If you do not specify the crypto map specific lifetime, the global value (or global default) is used.

See the [“Global Lifetime Values”](#) section on page 29-23 for more information on global lifetime values.

Send documentation comments to mdsfeedback-doc@cisco.com.

To set the SA lifetime for a specified crypto map entry, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# crypto map domain ipsec SampleMap 31 ips-hacl(config-crypto-map-ip)# | Place you in the crypto map configuration mode for the entry named SampleMap with 31 as its sequence number. |
| Step 3 | switch(config-crypto-map-ip)# set security-association lifetime seconds 8640 | Specifies a SA lifetime for this crypto map entry using different IPsec SA lifetimes than the global lifetimes, for the crypto map entry. |
| | switch(config-crypto-map-ip)# no set security-association lifetime seconds 8640 | Deletes the entry-specific configuration and reverts to the global settings. |
| Step 4 | switch(config-crypto-map-ip)# set security-association lifetime kilobytes 2560 | Configures the traffic-volume lifetime for this SA in kilobytes. The lifetime ranges from 2560 to 2147483647 kilobytes. |
| | switch(config-crypto-map-ip)# set security-association lifetime gigabytes 4000 | Configures the traffic-volume lifetime for this SA to time out after the specified amount of traffic (in gigabytes) have passed through the FCIP link using the SA. The lifetime ranges from 1 to 4095 gigabytes. |
| | switch(config-crypto-map-ip)# set security-association lifetime megabytes 5000 | Configures the traffic-volume lifetime for this SA in megabytes. The lifetime ranges from 3 to 4193280 megabytes. |
| | switch(config-crypto-map-ip)# no set security-association lifetime megabytes | Reverts to the global settings. |

The auto-peer Option


Setting peer address as **auto-peer** in the crypto map indicates that the destination endpoint of the traffic should be used as the peer address for the SA. Using the same crypto map, a unique SA can be setup to each of the endpoints in the subnet specified by the crypto map's ACL entry. Auto-peer simplifies configuration when traffic endpoints are IPsec capable. It is particularly useful for iSCSI, where the iSCSI hosts in the same subnet do not require separate configuration.

Figure 29-4 shows a scenario where the auto-peer option can simplify configuration. Using the auto-peer option, only one crypto map entry is needed for all the hosts from subnet X to setup SAs with the switch. Each host will setup its own SA, but will share the crypto map entry. Without the auto-peer option, each host needs one crypto map entry.

Refer to [Figure 29-6 on page 29-35](#) for more details.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 29-4 iSCSI with End-to-End IPsec Using the Auto-Peer Option



120879

To configure the auto-peer option, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# crypto map domain ipsec SampleMap 31 ips-hacl(config-crypto-map-ip)# | Place you in the crypto map configuration mode for the entry named SampleMap with 31 as its sequence number. |
| Step 3 | switch(config-crypto-map-ip)# set peer auto-peer | Directs the software to select (during the SA setup) the destination peer IP address dynamically. |
| | switch(config-crypto-map-ip)# no set peer auto-peer | Deletes the auto-peer configuration. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Perfect Forward Secrecy

To specify SA lifetime negotiation values, you can also optionally configure the perfect forward secrecy (PFS) value in the crypto map.

The PFS feature is disabled by default. If you set the PFS group, you can set one of DH groups: 1, 2, 5, or 14. If you do not specify a DH group, the software uses group 1 by default.

To configure the PFS value, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# crypto map domain ipsec SampleMap 31 ips-hacl(config-crypto-map-ip)# | Place you in the crypto map configuration mode for the entry named SampleMap with 31 as its sequence number. |
| Step 3 | switch(config-crypto-map-ip)# set pfs group 2 | Specifies that IPsec should ask for PFS when requesting new SAs for this crypto map entry, or should demand PFS in requests received from the IPsec peer. |
| | switch(config-crypto-map-ip)# no set pfs | Deletes the configured DH group and reverts to the factory default of disabling PFS. |

Crypto Map Set Interface Application

You need to apply a crypto map set to each interface through which IPsec traffic will flow. Applying the crypto map set to an interface instructs the switch to evaluate all the interface's traffic against the crypto map set and to use the specified policy during connection or SA negotiation on behalf of traffic to be protected by crypto.

You can apply only one crypto map set to an interface. You can apply the same crypto map to multiple interfaces. However, you cannot apply more than one crypto map set to each interface.

To apply a crypto map set to an interface, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# interface gigabitethernet 4/1 switch(config-if)# | Selects the required Gigabit Ethernet interface (and subinterface, if required) to which the IPsec crypto map is to be applied. |
| Step 3 | switch(config-if)# crypto map domain ipsec cm10 | Applies the crypto map set to the selected interface. |
| Step 4 | switch(config-if)# no crypto map domain ipsec | Deletes the crypto map that is currently applied to this interface. |

Send documentation comments to mdsfeedback-doc@cisco.com.

IPsec Maintenance

Certain configuration changes will only take effect when negotiating subsequent security associations. If you want the new settings to take immediate effect, you must clear the existing security associations so that they will be re-established with the changed configuration. If the switch is actively processing IPsec traffic, it is desirable to clear only the portion of the security association database that would be affected by the configuration changes (that is, clear only the security associations established by a given crypto map set). Clearing the full security association database should be reserved for large-scale changes, or when the router is processing very little other IPsec traffic.



Caution

Using the **clear crypto sa** command without parameters will clear out the full SA database, which will clear out active security sessions. You may also specify the peer, map, or entry keywords to clear out only a subset of the SAD.

Use the **clear crypto sa** command to clear all or part of the SAD.

```
switch# clear crypto sa domain ipsec interface gigabitethernet 2/1 inbound sa 1
```



Tip

You can obtain the SA index from the output of the **show crypto sa domain interface gigabitethernet slot/port** command.

Global Lifetime Values

If you have not configured a lifetime in the crypto map entry, the global lifetime values are used when negotiating new IPsec SAs.

You can configure two lifetimes: timed or traffic-volume. A SA expires after the first of these lifetimes is reached. The default lifetimes are 3,600 seconds (one hour) and 450 GB.

If you change a global lifetime, the new lifetime value will not be applied to currently existing SAs, but will be used in the negotiation of subsequently established SAs. If you wish to use the new values immediately, you can clear all or part of the SA database.

Assuming that the particular crypto map entry does not have lifetime values configured, when the switch requests new SAs it will specify its global lifetime values in the request to the peer; it will use this value as the lifetime of the new SAs. When the switch receives a negotiation request from the peer, it uses the value determined by the IKE version in use:

- If you use IKEv1 to setup IPsec SAs, the SA lifetime values are chosen to be the smaller of the two proposals. The same values are programmed on both the ends of the tunnel.
- If you use IKEv2 to setup IPsec SAs, SAs on each end has its own set up of lifetime values and thus the SAs on both sides expire independently.

The SA (and corresponding keys) will expire according to whichever comes sooner, either after the specified amount of time (in seconds) has passed or after the specified amount of traffic (in bytes) has passed.

A new SA is negotiated before the lifetime threshold of the existing SA is reached, to ensure that negotiation completes before the existing SA expires.

Send documentation comments to mdsfeedback-doc@cisco.com.

The new SA is negotiated when one of the following thresholds is reached (whichever comes first):

- 30 seconds before the lifetime expires or
- Approximately 10% of the lifetime in bytes remain

If no traffic has passed through when the lifetime expires, a new SA is not negotiated. Instead, a new SA will be negotiated only when IPsec sees another packet that should be protected.

To configure global SA lifetimes, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# crypto global domain ipsec security-association lifetime seconds 86400 | Configures the global timed lifetime for IPsec SAs to time out after the specified number of seconds have passed. The global lifetime ranges from 120 to 86400 seconds. |
| | switch(config)# no crypto global domain ipsec security-association lifetime seconds 86400 | Reverts to the factory default of 3,600 seconds. |
| Step 3 | switch(config)# crypto global domain ipsec security-association lifetime gigabytes 4000 | Configures the global traffic-volume lifetime for IPsec SAs to time out after the specified amount of traffic (in gigabytes) have passed through the FCIP link using the SA. The global lifetime ranges from 1 to 4095 gigabytes. |
| | switch(config)# crypto global domain ipsec security-association lifetime kilobytes 2560 | Configures the global traffic-volume lifetime in kilobytes. The global lifetime ranges from 2560 to 2147483647 kilobytes. |
| | switch(config)# crypto global domain ipsec security-association lifetime megabytes 5000 | Configures the global traffic-volume lifetime in megabytes. The global lifetime ranges from 3 to 4193280 megabytes. |
| | switch(config)# no crypto global domain ipsec security-association lifetime megabytes | Reverts to the factory default of 450 GB regardless of what value is currently configured. |

Displaying IKE Configurations

You can verify the IKE information by using the **show** set of commands. See Examples 29-1 to 29-5.

Example 29-1 *Displays the Parameters Configured for Each IKE policy*

```
switch# show crypto ike domain ipsec
keepalive 60000
```

Example 29-2 *Displays the Initiator Configuration*

```
switch# show crypto ike domain ipsec initiator
initiator version 1 address 1.1.1.1
initiator version 1 address 1.1.1.2
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 29-3 Displays the Key Configuration

```
switch# show crypto ike domain ipsec key
key abcdefgh address 1.1.1.1
key bcdefghi address 1.1.2.1
```

Example 29-4 Displays the Currently Established Policies for IKE

```
switch# show crypto ike domain ipsec policy 1
Priority 1, auth pre-shared, lifetime 6000 secs, encryption 3des, hash md5, DH group 5
Priority 3, auth pre-shared, lifetime 86300 secs, encryption aes, hash sha1, DH group 1
```

Example 29-5 Displays the Currently Established SAs for IKE

```
switch# show crypto ike domain ipsec sa
Tunn    Local Addr          Remote Addr          Encr    Hash    Auth Method    Lifetime
-----
1*      172.22.31.165[500]    172.22.31.166[500]  3des    sha1    preshared key  86400
2       172.22.91.174[500]    172.22.91.173[500]  3des    sha1    preshared key  86400
-----
NOTE: tunnel id ended with * indicates an IKEv1 tunnel
```

Displaying IPsec Configurations

You can verify the IPsec information by using the **show** set of commands. See Examples 29-6 to 29-20.

Example 29-6 Displays IP ACL Information

```
switch# show ip access-list usage
Access List Name/Number    Filters IF    Status    Creation Time
-----
acl110                     1            0        active    Mon Mar  2 05:07:20 1981
acl1100                    1            0        active    Mon Mar  2 05:07:20 1981
acl1100subnet              1            0        active    Mon Mar  2 05:07:20 1981
```

Example 29-7 Displays Information for the Specified ACL

```
switch# show ip access-list acl110
ip access-list acl110 permit ip 10.10.10.0 0.0.0.255 10.10.10.0 0.0.0.255 (0 matches)
```

In [Example 29-7](#), the display output match is only displayed of an interface (not the crypto map) meets this criteria.

Example 29-8 Displays the Transform Set Configuration

```
switch# show crypto transform-set domain ipsec
Transform set: 3des-md5 {esp-3des esp-md5-hmac}
will negotiate {tunnel}
Transform set: des-md5 {esp-des esp-md5-hmac}
will negotiate {tunnel}
Transform set: test {esp-aes-128-cbc esp-md5-hmac}
will negotiate {tunnel}
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 29-9 Displays All Configured Crypto Maps

```
switch# show crypto map domain ipsec
Crypto Map "cm10" 1 ipsec
  Peer = Auto Peer
  IP ACL = acl10
    permit ip 10.10.10.0 255.255.255.0 10.10.10.0 255.255.255.0
  Transform-sets: 3des-md5, des-md5,
  Security Association Lifetime: 4500 megabytes/3600 seconds
  PFS (Y/N): N
  Interface using crypto map set cm10:
    GigabitEthernet4/1
Crypto Map "cm100" 1 ipsec
  Peer = Auto Peer
  IP ACL = acl100
    permit ip 10.10.100.0 255.255.255.0 10.10.100.0 255.255.255.0
  Transform-sets: 3des-md5, des-md5,
  Security Association Lifetime: 4500 megabytes/3600 seconds
  PFS (Y/N): N
  Interface using crypto map set cm100:
    GigabitEthernet4/2
```

Example 29-10 Displays the Crypto Map Information for a Specific Interface

```
switch# show crypto map domain ipsec interface gigabitethernet 4/1
Crypto Map "cm10" 1 ipsec
  Peer = Auto Peer
  IP ACL = acl10
    permit ip 10.10.10.0 255.255.255.0 10.10.10.0 255.255.255.0
  Transform-sets: 3des-md5, des-md5,
  Security Association Lifetime: 4500 megabytes/120 seconds
  PFS (Y/N): N
  Interface using crypto map set cm10:
    GigabitEthernet4/1
```

Example 29-11 Displays the Specified Crypto Map Information

```
switch# show crypto map domain ipsec tag cm100
Crypto Map "cm100" 1 ipsec
  Peer = Auto Peer
  IP ACL = acl100
    permit ip 10.10.100.0 255.255.255.0 10.10.100.0 255.255.255.0
  Transform-sets: 3des-md5, des-md5,
  Security Association Lifetime: 4500 megabytes/120 seconds
  PFS (Y/N): N
  Interface using crypto map set cm100:
    GigabitEthernet4/2
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 29-12 Displays SA Association for the Specified Interface

```
switch# show crypto sad domain ipsec interface gigabitethernet 4/1
interface: GigabitEthernet4/1
  Crypto map tag: cm10, local addr. 10.10.10.1
  protected network:
    local ident (addr/mask): (10.10.10.0/255.255.255.0)
    remote ident (addr/mask): (10.10.10.4/255.255.255.255)
    current_peer: 10.10.10.4
      local crypto endpt.: 10.10.10.1, remote crypto endpt.: 10.10.10.4
      mode: tunnel, crypto algo: esp-3des, auth algo: esp-md5-hmac
      current outbound spi: 0x30e000f (51249167), index: 0
      lifetimes in seconds:: 120
      lifetimes in bytes:: 423624704
      current inbound spi: 0x30e0000 (51249152), index: 0
      lifetimes in seconds:: 120
      lifetimes in bytes:: 423624704
```

Example 29-13 Displays All SA Associations

```
switch# show crypto sad domain ipsec
interface: GigabitEthernet4/1
  Crypto map tag: cm10, local addr. 10.10.10.1
  protected network:
    local ident (addr/mask): (10.10.10.0/255.255.255.0)
    remote ident (addr/mask): (10.10.10.4/255.255.255.255)
    current_peer: 10.10.10.4
      local crypto endpt.: 10.10.10.1, remote crypto endpt.: 10.10.10.4
      mode: tunnel, crypto algo: esp-3des, auth algo: esp-md5-hmac
      current outbound spi: 0x30e000f (51249167), index: 0
      lifetimes in seconds:: 120
      lifetimes in bytes:: 423624704
      current inbound spi: 0x30e0000 (51249152), index: 0
      lifetimes in seconds:: 120
      lifetimes in bytes:: 423624704
```

Example 29-14 Displays Information About the Policy Database

```
switch# show crypto spd domain ipsec
Policy Database for interface: GigabitEthernet4/1, direction: Both
# 0:      deny  udp any port eq 500 any
# 1:      deny  udp any any port eq 500
# 2:      permit ip 10.10.10.0 255.255.255.0 10.10.10.0 255.255.255.0
# 63:     deny  ip any any
Policy Database for interface: GigabitEthernet4/2, direction: Both
# 0:      deny  udp any port eq 500 any <-----UDP default entry
# 1:      deny  udp any any port eq 500 <-----UDP default entry
# 3:      permit ip 10.10.100.0 255.255.255.0 10.10.100.0 255.255.255.0
# 63:     deny  ip any any <-----Clear text default entry
```

Example 29-15 Displays SPD Information for a Specific Interface

```
switch# show crypto spd domain ipsec interface gigabitethernet 4/2
Policy Database for interface: GigabitEthernet3/1, direction: Both
# 0:      deny  udp any port eq 500 any
# 1:      deny  udp any any port eq 500
# 2:      permit ip 10.10.10.0 255.255.255.0 10.10.10.0 255.255.255.0
# 127:     deny  ip any any
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 29-16 Displays Detailed iSCSI Session Information for a Specific Interface

```
switch# show iscsi session detail
Initiator ign.1987-05.com.cisco:01.9f39f09c7468 (ips-host16.cisco.com)
  Initiator ip addr (s): 10.10.10.5
  Session #1 (index 24)
    Discovery session, ISID 00023d000001, Status active

  Session #2 (index 25)
    Target ibml
    VSAN 1, ISID 00023d000001, TSIH 0, Status active, no reservation
    Type Normal, ExpCmdSN 42, MaxCmdSN 57, Barrier 0
    MaxBurstSize 0, MaxConn 1, DataPDUInOrder Yes
    DataSeqInOrder Yes, InitialR2T Yes, ImmediateData No
    Registered LUN 0, Mapped LUN 0
    Stats:
      PDU: Command: 41, Response: 41
      Bytes: TX: 21388, RX: 0
    Number of connection: 1
    Connection #1
      iSCSI session is protected by IPsec <-----The iSCSI session protection status
      Local IP address: 10.10.10.4, Peer IP address: 10.10.10.5
      CID 0, State: Full-Feature
      StatSN 43, ExpStatSN 0
      MaxRecvDSLength 131072, our_MaxRecvDSLength 262144
      CSG 3, NSG 3, min_pdu_size 48 (w/ data 48)
      AuthMethod none, HeaderDigest None (len 0), DataDigest None (len 0)
      Version Min: 0, Max: 0
      FC target: Up, Reorder PDU: No, Marker send: No (int 0)
      Received MaxRecvDSLen key: Yes
```

Example 29-17 Displays FCIP Information for a Specific Interface

```
switch# show interface fcip 1
fcip1 is trunking
  Hardware is GigabitEthernet
  Port WWN is 20:50:00:0d:ec:08:6c:c0
  Peer port WWN is 20:10:00:05:30:00:a7:9e
  Admin port mode is auto, trunk mode is on
  Port mode is TE
  Port vsan is 1
  Speed is 1 Gbps
  Trunk vsans (admin allowed and active) (1)
  Trunk vsans (up) (1)
  Trunk vsans (isolated) ( )
  Trunk vsans (initializing) ( )
  Using Profile id 1 (interface GigabitEthernet2/1)
  Peer Information
    Peer Internet address is 10.10.11.1 and port is 3225
    FCIP tunnel is protected by IPsec <-----The FCIP tunnel protection status
    Write acceleration mode is off
    Tape acceleration mode is off
    Tape Accelerator flow control buffer size is 256 KBytes
    IP Compression is disabled
    Special Frame is disabled
    Maximum number of TCP connections is 2
    Time Stamp is disabled
    QOS control code point is 0
    QOS data code point is 0
    B-port mode disabled
  TCP Connection Information
    2 Active TCP connections
```


Send documentation comments to mdsfeedback-doc@cisco.com.

```
Control connection: Local 10.10.11.2:3225, Remote 10.10.11.1:65520
Data connection: Local 10.10.11.2:3225, Remote 10.10.11.1:65522
2 Attempts for active connections, 0 close of connections
TCP Parameters
Path MTU 1400 bytes
Current retransmission timeout is 200 ms
Round trip time: Smoothed 2 ms, Variance: 1
Advertized window: Current: 124 KB, Maximum: 124 KB, Scale: 6
Peer receive window: Current: 123 KB, Maximum: 123 KB, Scale: 6
Congestion window: Current: 53 KB, Slow start threshold: 48 KB
Current Send Buffer Size: 124 KB, Requested Send Buffer Size: 0 KB
CWM Burst Size: 50 KB
5 minutes input rate 128138888 bits/sec, 16017361 bytes/sec, 7937 frames/sec
5 minutes output rate 179275536 bits/sec, 22409442 bytes/sec, 46481 frames/sec
10457037 frames input, 21095415496 bytes
308 Class F frames input, 32920 bytes
10456729 Class 2/3 frames input, 21095382576 bytes
9907495 Reass frames
0 Error frames timestamp error 0
63792101 frames output, 30250403864 bytes
472 Class F frames output, 46816 bytes
63791629 Class 2/3 frames output, 30250357048 bytes
0 Error frames
```

Example 29-18 Displays the Global IPsec Statistics for the Switch

```
switch# show crypto global domain ipsec
IPSec global statistics:
  Number of crypto map sets: 3
  IKE transaction stats: 0 num, 256 max
  Inbound SA stats: 0 num
  Outbound SA stats: 0 num
```

Example 29-19 Displays the IPsec Statistics for the Specified Interface

```
switch# show crypto global domain ipsec interface gigabitethernet 3/1
IPSec interface statistics:
  IKE transaction stats: 0 num
  Inbound SA stats: 0 num, 512 max
  Outbound SA stats: 0 num, 512 max
```

Example 29-20 Displays the Global SA Lifetime Values


```
switch# show crypto global domain ipsec security-association lifetime
Security Association Lifetime: 450 gigabytes/3600 seconds
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Sample FCIP Configuration

Figure 29-5 focuses on implementing IPsec for one FCIP link (Tunnel 2). Tunnel 2 carries encrypted data between MDS A and MDS C.

Figure 29-5 IP Security Usage in an FCIP Scenario



To configure IPsec for the FCIP scenario shown in Figure 29-5, follow these steps:

Step 1 Enable IKE and IPsec in Switch MDS A.

```
sw10.1.1.100# conf t
sw10.1.1.100(config)# crypto ike enable
sw10.1.1.100(config)# crypto ipsec enable
```

Step 2 Configure IKE in Switch MDS A.

```
sw10.1.1.100(config)# crypto ike domain ipsec
sw10.1.1.100(config-ike-ipsec)# key ctct address 10.10.100.232
sw10.1.1.100(config-ike-ipsec)# policy 1
sw10.1.1.100(config-ike-ipsec-policy)# encryption 3des
sw10.1.1.100(config-ike-ipsec-policy)# hash md5
sw10.1.1.100(config-ike-ipsec-policy)# end
sw10.1.1.100#
```

Step 3 Configure the ACLs in Switch MDS A.

```
sw10.1.1.100# conf t
sw10.1.1.100(config)# ip access-list acl1 permit ip 10.10.100.231 0.0.0.0 10.10.100.232 0.0.0.0
```

Step 4 Configure the transform set in Switch MDS A.

```
sw10.1.1.100(config)# crypto transform-set domain ipsec tfs-02 esp-aes 128 esp-sha1-hmac
```

Step 5 Configure the crypto map in Switch MDS A.

```
sw10.1.1.100(config)# crypto map domain ipsec cmap-01 1
sw10.1.1.100(config-crypto-map-ip)# match address acl1
sw10.1.1.100(config-crypto-map-ip)# set peer 10.10.100.232
sw10.1.1.100(config-crypto-map-ip)# set transform-set tfs-02
sw10.1.1.100(config-crypto-map-ip)# set security-association lifetime seconds 120
sw10.1.1.100(config-crypto-map-ip)# set security-association lifetime gigabytes 3000
sw10.1.1.100(config-crypto-map-ip)# set pfs group5
sw10.1.1.100(config-crypto-map-ip)# end
sw10.1.1.100#
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Step 6 Bind the interface to the crypto map set in Switch MDS A.

```
sw10.1.1.100# conf t
sw10.1.1.100(config)# int gigabitethernet 7/1
sw10.1.1.100(config-if)# ip addr 10.10.100.231 255.255.255.0
sw10.1.1.100(config-if)# crypto map domain ipsec cmap-01
sw10.1.1.100(config-if)# no shut
sw10.1.1.100(config-if)# exit
sw10.1.1.100(config)#
```

Step 7 Configure FCIP in Switch MDS A.

```
sw10.1.1.100(config)# fcip enable
sw10.1.1.100(config)# fcip profile 2
sw10.1.1.100(config-profile)# ip address 10.10.100.231
sw10.1.1.100(config-profile)# int fcip 2
sw10.1.1.100(config-if)# peer-info ipaddr 10.10.100.232
sw10.1.1.100(config-if)# use-profile 2
sw10.1.1.100(config-if)# no shut
sw10.1.1.100(config-if)# end
sw10.1.1.100#
```

Step 8 Verify the configuration in Switch MDS A.

```
sw10.1.1.100# show crypto global domain ipsec security-association lifetime
Security Association Lifetime: 4500 megabytes/3600 seconds

sw10.1.1.100# show crypto map domain ipsec
Crypto Map "cmap-01" 1 ipsec
  Peer = 10.10.100.232
  IP ACL = acl1
    permit ip 10.10.100.231 255.255.255.255 10.10.100.232 255.255.255.255
  Transform-sets: tfs-02,
  Security Association Lifetime: 3000 gigabytes/120 seconds
  PFS (Y/N): Y
  PFS Group: group5
Interface using crypto map set cmap-01:
  GigabitEthernet7/1

sw10.1.1.100# show crypto transform-set domain ipsec
Transform set: tfs-02 {esp-aes 128 esp-shal-hmac}
  will negotiate {tunnel}

sw10.1.1.100# show crypto spd domain ipsec
Policy Database for interface: GigabitEthernet7/1, direction: Both
# 0:      deny  udp any port eq 500 any
# 1:      deny  udp any any port eq 500
# 2:      permit ip 10.10.100.231 255.255.255.255 10.10.100.232 255.255.255.255
# 63:     deny  ip any any

sw10.1.1.100# show crypto ike domain ipsec
keepalive 3600

sw10.1.1.100# show crypto ike domain ipsec key
key ctct address 10.10.100.232

sw10.1.1.100# show crypto ike domain ipsec policy
Priority 1, auth pre-shared, lifetime 86300 secs, encryption 3des, hash md5, DH group 1
```

Step 9 Enable IKE and IPsec in Switch MDS C.

```
sw11.1.1.100# conf t
sw11.1.1.100(config)# crypto ike enable
sw11.1.1.100(config)# crypto ipsec enable
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Step 10 Configure IKE in Switch MDS C.

```
sw11.1.1.100(config)# crypto ike domain ipsec
sw11.1.1.100(config-ike-ipsec)# key ctct address 10.10.100.231
sw11.1.1.100(config-ike-ipsec)# policy 1
sw11.1.1.100(config-ike-ipsec-policy)# encryption 3des
sw11.1.1.100(config-ike-ipsec-policy)# hash md5
sw11.1.1.100(config-ike-ipsec-policy)# exit
sw11.1.1.100(config-ike-ipsec)# end
sw11.1.1.100#
```

Step 11 Configure the ACLs in Switch MDS C.

```
sw11.1.1.100# conf t
sw11.1.1.100(config)# ip access-list acl1 permit ip 10.10.100.232 0.0.0.0 10.10.100.231
0.0.0.0
```

Step 12 Configure the transform set in Switch MDS C.

```
sw11.1.1.100(config)# crypto transform-set domain ipsec tfs-02 esp-aes 128 esp-sha1-hmac
```

Step 13 Configure the crypto map in Switch MDS C.

```
sw11.1.1.100(config)# crypto map domain ipsec cmap-01 1
sw11.1.1.100(config-crypto-map-ip)# match address acl1
sw11.1.1.100(config-crypto-map-ip)# set peer 10.10.100.231
sw11.1.1.100(config-crypto-map-ip)# set transform-set tfs-02
sw11.1.1.100(config-crypto-map-ip)# set security-association lifetime seconds 120
sw11.1.1.100(config-crypto-map-ip)# set security-association lifetime gigabytes 3000
sw11.1.1.100(config-crypto-map-ip)# set pfs group5
sw11.1.1.100(config-crypto-map-ip)# exit
sw11.1.1.100(config)#
```

Step 14 Bind the interface to the crypto map set in Switch MDS C.

```
sw11.1.1.100(config)# int gigabitethernet 1/2
sw11.1.1.100(config-if)# ip addr 10.10.100.232 255.255.255.0
sw11.1.1.100(config-if)# crypto map domain ipsec cmap-01
sw11.1.1.100(config-if)# no shut
sw11.1.1.100(config-if)# exit
sw11.1.1.100(config)#
```

Step 15 Configure FCIP in Switch MDS C.

```
sw11.1.1.100(config)# fcip enable
sw11.1.1.100(config)# fcip profile 2
sw11.1.1.100(config-profile)# ip address 10.10.100.232
sw11.1.1.100(config-profile)# int fcip 2
sw11.1.1.100(config-if)# peer-info ipaddr 10.10.100.231
sw11.1.1.100(config-if)# use-profile 2
sw11.1.1.100(config-if)# no shut
sw11.1.1.100(config-if)# exit
sw11.1.1.100(config)# exit
```

Step 16 Verify the configuration in Switch MDS C.

```
sw11.1.1.100# show crypto global domain ipsec security-association lifetime
Security Association Lifetime: 4500 megabytes/3600 seconds
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
sw11.1.1.100# show crypto map domain ipsec
Crypto Map "cmap-01" 1 ipsec
  Peer = 10.10.100.231
  IP ACL = acl1
    permit ip 10.10.100.232 255.255.255.255 10.10.100.231 255.255.255.255
  Transform-sets: tfs-02,
  Security Association Lifetime: 3000 gigabytes/120 seconds
  PFS (Y/N): Y
    PFS Group: group5
Interface using crypto map set cmap-01:
  GigabitEthernet1/2
```

```
sw11.1.1.100# show crypto spd domain ipsec
Policy Database for interface: GigabitEthernet1/2, direction: Both
# 0:      deny  udp any port eq 500 any
# 1:      deny  udp any any port eq 500
# 2:      permit ip 10.10.100.232 255.255.255.255 10.10.100.231 255.255.255.255
# 63:     deny  ip any any
```

```
sw11.1.1.100# show crypto sad domain ipsec
interface: GigabitEthernet1/2
  Crypto map tag: cmap-01, local addr. 10.10.100.232
  protected network:
    local ident (addr/mask): (10.10.100.232/255.255.255.255)
    remote ident (addr/mask): (10.10.100.231/255.255.255.255)
    current_peer: 10.10.100.231
      local crypto endpt.: 10.10.100.232, remote crypto endpt.: 10.10.100.231
      mode: tunnel, crypto algo: esp-3des, auth algo: esp-md5-hmac
      current outbound spi: 0x38f96001 (955867137), index: 29
      lifetimes in seconds:: 120
      lifetimes in bytes:: 3221225472000
      current inbound spi: 0x900b011 (151040017), index: 16
      lifetimes in seconds:: 120
      lifetimes in bytes:: 3221225472000
```

```
sw11.1.1.100# show crypto transform-set domain ipsec
Transform set: tfs-02 {esp-aes 128 esp-sha1-hmac}
  will negotiate {tunnel}
```

```
sw11.1.1.100# show crypto ike domain ipsec
keepalive 3600
```

```
sw11.1.1.100# show crypto ike domain ipsec key
```

```
key ctct address 10.10.100.231
```

```
sw11.1.1.100# show crypto ike domain ipsec policy
Priority 1, auth pre-shared, lifetime 86300 secs, encryption 3des, hash md5, DH
group 1
```

```
sw11.1.1.100# show crypto ike domain ipsec sa
```

| Tunn | Local Addr | Remote Addr | Encr | Hash | Auth Method | Lifetime |
|------|--------------------|--------------------|------|------|---------------|----------|
| 1* | 10.10.100.232[500] | 10.10.100.231[500] | 3des | md5 | preshared key | 86300 |

NOTE: tunnel id ended with * indicates an IKEv1 tunnel

Send documentation comments to mdsfeedback-doc@cisco.com.

Step 17 Verify the configuration in Switch MDS A.

```
sw10.1.1.100# show crypto sad domain ipsec
interface: GigabitEthernet7/1
  Crypto map tag: cmap-01, local addr. 10.10.100.231
  protected network:
    local ident (addr/mask): (10.10.100.231/255.255.255.255)
    remote ident (addr/mask): (10.10.100.232/255.255.255.255)
    current_peer: 10.10.100.232
      local crypto endpt.: 10.10.100.231, remote crypto endpt.: 10.10.100.232
      mode: tunnel, crypto algo: esp-3des, auth algo: esp-md5-hmac
      current outbound spi: 0x900b01e (151040030), index: 10
      lifetimes in seconds:: 120
      lifetimes in bytes:: 3221225472000
      current inbound spi: 0x38fe700e (956198926), index: 13
      lifetimes in seconds:: 120
      lifetimes in bytes:: 3221225472000

sw10.1.1.100# show crypto ike domain ipsec sa
Tunn Local Addr      Remote Addr      Encr  Hash  Auth Method      Lifetime
-----
  1 10.10.100.231[500] 10.10.100.232[500] 3des  md5   preshared key      86300
```


You have now configured IPsec in both switches MDS A and MDS C.

Sample iSCSI Configuration

[Figure 29-6](#) focuses on the iSCSI session between MDS A and the hosts in subnet 12.12.1/24. Using the **auto-peer** option, when any host from the subnet 12.12.1.0/24 tries to connect to MDS's Gigabit Ethernet port 7/1, an SA is created between the hosts and MDS. With auto-peer, only one crypto map is necessary to create SAs for all the hosts in the same subnet. Without auto-peer, you need one crypto map entry per host.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 29-6 iSCSI with End-to-End IPsec



To configure IPsec for the iSCSI scenario shown in [Figure 29-6](#), follow these steps:

Step 1 Configure the ACLs in Switch MDS A.

```
sw10.1.1.100# conf t
sw10.1.1.100(config)# ip access-list acl1 permit ip 10.10.1.0 0.0.0.255 12.12.1.0
0.0.0.255
```

Step 2 Configure the transform set in Switch MDS A.

```
sw10.1.1.100(config)# crypto transform-set domain ipsec tfs-01 esp-3des esp-md5-hmac
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Step 3 Configure the crypto map in Switch MDS A.

```
sw10.1.1.100(config)# crypto map domain ipsec cmap-01 1
sw10.1.1.100(config-crypto-map-ip)# match address acl1
sw10.1.1.100(config-crypto-map-ip)# set peer auto-peer
sw10.1.1.100(config-crypto-map-ip)# set transform-set tfs-01
sw10.1.1.100(config-crypto-map-ip)# end
sw10.1.1.100#
```

Step 4 Bind the interface to the crypto map set in Switch MDS A.

```
sw10.1.1.100# conf t
sw10.1.1.100(config)# int gigabitethernet 7/1
sw10.1.1.100(config-if)# ip addr 10.10.1.123 255.255.255.0
sw10.1.1.100(config-if)# crypto map domain ipsec cmap-01
sw10.1.1.100(config-if)# no shut
sw10.1.1.100(config-if)# end
sw10.1.1.100#
```

You have now configured IPsec in MDS A using the Cisco MDS IPsec and iSCSI features.

Default Settings

Table 29-3 lists the default settings for IKE parameters.

Table 29-3 **Default IKE Parameters**

| Parameters | Default |
|---------------------------------------|---|
| IKE | Disabled. |
| IKE version | IKE version 2. |
| IKE encryption algorithm | 3DES. |
| IKE hash algorithm | SHA. |
| IKE authentication method | Not configurable (uses preshared keys). |
| IKE DH group identifier | Group 1. |
| IKE lifetime association | 86,400 00 seconds (equals 24 hours). |
| IKE keepalive time for each peer (v2) | 3,600 seconds (equals one hour). |

Table 29-4 lists the default settings for IPsec parameters.

Table 29-4 **Default IPsec Parameters**

| Parameters | Default |
|--|---------------------------|
| IPsec | Disabled. |
| Applying IPsec to the traffic. | Deny—allowing clear text. |
| IPsec PFS | Disabled. |
| IPsec global lifetime (traffic-volume) | 450 Gigabytes. |
| IPsec global lifetime (time) | 3,600 seconds (one hour). |



CHAPTER 30

Configuring Call Home

Call Home provides e-mail-based notification of critical system events. A versatile range of message formats are available for optimal compatibility with pager services, standard e-mail, or XML-based automated parsing applications. Common uses of this feature may include direct paging of a network support engineer, e-mail notification to a Network Operations Center, and utilization of Cisco AutoNotify services for direct case generation with the Technical Assistance Center.

As of Cisco SAN-OS Release 2.0(1b), the Call Home feature provides message throttling capabilities. Periodic inventory messages, port syslog messages and RMON alert messages are added to the list of deliverable Call Home messages. If required you can also use the Cisco Fabric Services application to distribute the Call Home configuration to all other switches in the fabric.

This chapter includes the following sections:

- [Call Home Features, page 30-2](#)
- [Call Home Configuration Process, page 30-3](#)
- [Cisco AutoNotify, page 30-2](#)
- [Call Home Configuration Process, page 30-3](#)
- [Destination Profiles, page 30-4](#)
- [Alert Groups, page 30-6](#)
- [Call Home Message Levels, page 30-8](#)
- [Syslog-based Alerts, page 30-9](#)
- [RMON-based Alerts, page 30-9](#)
- [E-Mail Options, page 30-10](#)
- [Periodic Inventory Notification, page 30-11](#)
- [Duplicate Message Throttle, page 30-11](#)
- [Call Home Enable Function, page 30-12](#)
- [Call Home Configuration Distribution, page 30-12](#)
- [Call Home Communications Test, page 30-14](#)
- [Displaying Call Home Information, page 30-14](#)
- [Default Settings, page 30-18](#)
- [Event Triggers, page 30-19](#)
- [Call Home Message Levels, page 30-20](#)
- [Message Contents, page 30-21](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

Call Home Features

The Call Home functionality is available directly through the Cisco MDS 9000 Family. It provides multiple Call Home profiles (also referred to as Call Home destination profiles), each with separate potential destinations. You can define your own destination-profiles in addition to predefined profiles.

The Call Home function can even leverage support from Cisco Systems or another support partner. Flexible message delivery and format options make it easy to integrate specific support requirements.

The Call Home feature offers the following advantages:

- Fixed set of predefined alerts and trigger events on the switch.
- Automatic execution and attachment of relevant command output.
- Multiple message format options:
 - Short Text—Suitable for pagers or printed reports.
 - Plain Text—Full formatted message information suitable for human reading.
 - XML—Matching readable format using Extensible Markup Language (XML) and document type definitions (DTDs) named Messaging Markup Language (MML). The MML DTD is published on the Cisco.com website at <http://www.cisco.com/>. The XML format enables communication with the Cisco Systems Technical Assistance Center.
- Multiple concurrent message destinations. You can configure up to 50 e-mail destination addresses for each destination-profile.
- Multiple message categories including system, environment, switching module hardware, supervisor module, hardware, inventory, syslog, RMON, and test.

Cisco AutoNotify

For those who have service contracts directly with Cisco Systems, automatic case generation with the Technical Assistance Center is possible by registering with the AutoNotify service. AutoNotify provides fast time to resolution of system problems by providing a direct notification path to Cisco customer support.

The AutoNotify feature requires several Call Home parameters to be configured, including certain contact information, e-mail server, and an XML destination profile as specified in the Service Activation document found on the Cisco.com web site at:

http://www.cisco.com/univercd/cc/td/doc/product/voice/c_callmg/3_3/service/serv332/ccmsrvs/sssrvact.htm

To configure a Cisco MDS 9000 Family switch to use the AutoNotify service, an XML destination profile must be configured to send messages to Cisco. Specific setup, activation, and e-mail address information is found on the Cisco.com web site at:

http://www.cisco.com/warp/customer/cc/serv/mkt/sup/tsssv/opmsup/smtan/anoti_ds.htm

To register, the following items are required:

- The SMARTnet contract number covering your Cisco MDS 9000 Family switch.
- Your name, company address, your e-mail address, and your Cisco.com ID.
- The exact product number of your Cisco MDS 9000 Family switch. For example, some valid product numbers include: DS-C6509 and DS-C9216-K9.
- The serial number of your Cisco MDS 9000 Family switch. This can be obtained by looking at the serial number label on the back of the switch (next to the power supply).

Send documentation comments to mdsfeedback-doc@cisco.com.

The ContractID, CustomerID, SiteID, and SwitchPriority parameters are not required by the AutoNotify feature. They are only intended to be used as additional information by Cisco customers and service partners.

se the **show sprom backplane 1** command or the **show license host-id** command to obtain the switch serial number.

Call Home Configuration Process

The actual configuration of Call Home depends on how you intend to use the feature. Some points to consider include:

- An e-mail server and at least one destination profile (predefined or user-defined) must be configured. The destination profile(s) used depends on whether the receiving entity is a pager, e-mail, or automated service such as Cisco AutoNotify.
- The contact name (SNMP server contact), phone, and street address information must be configured before Call Home is enabled. This is required to determine the origin of messages received.
- The Cisco MDS 9000 switch must have IP connectivity to an e-mail server.
- If Cisco AutoNotify is used, an active service contract must cover the device being configured.

To configure Call Home, follow these steps:

-
- | | |
|---------------|---|
| Step 1 | Assign contact information. |
| Step 2 | Configure destination profiles. |
| Step 3 | Associate one or more alert groups to each profile as required by your network. |
| Step 4 | Enable or disable Call Home. |
| Step 5 | Test Call Home messages. |
-

Contact Information

It is mandatory for each switch to include e-mail, phone, and street address information. It is optional to include the contract ID, customer ID, site ID, and switch priority information.

To assign the contact information, follow these steps:

| | Command | Purpose |
|---------------|--|---|
| Step 1 | switch# confi g t | Enters configuration mode. |
| Step 2 | switch# snmp-server contact personname@companyname.com | Configures the SNMP contact name. |
| Step 3 | switch(config)# callhome switch(config-callhome) # | Enters the Call Home submode. |
| Step 4 | switch(config-callhome)# email-contact username@company.com | Assigns the customer's e-mail address. Up to 128 alphanumeric characters are accepted in e-mail address format. |
| | | Note You can use any valid e-mail address. You cannot use spaces. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|---------|--|--|
| Step 5 | <code>switch(config-callhome) # phone-contact +1-800-123-4567</code> | Assigns the customer's phone number. Up to 20 alphanumeric characters are accepted in international format. Note You cannot use spaces. Be sure to use the + prefix before the number. |
| Step 6 | <code>switch(config-callhome) # streetaddress 1234 Picaboo Street, Any city, Any state, 12345</code> | Assigns the customer's street address where the equipment is located. Up to 256 alphanumeric characters are accepted in free format. |
| Step 7 | <code>switch(config-callhome) # switch-priority 0</code> | Assigns the switch priority, with 0 being the highest priority and 7 the lowest. Tip Use this field to create a hierarchical management structure. |
| Step 8 | <code>switch(config-callhome) # customer-id Customer1234</code> | Optional. Identifies the customer ID. Up to 256 alphanumeric characters are accepted in free format. |
| Step 9 | <code>switch(config-callhome) # site-id Site1ManhattanNY</code> | Optional. Identifies the customer site ID. Up to 256 alphanumeric characters are accepted in free format. |
| Step 10 | <code>switch(config-callhome) # contract-id Company1234</code> | Assigns the customer ID for the switch. Up to 64 alphanumeric characters are accepted in free format. |

Destination Profiles

A destination profile contains the required delivery information for an alert notification. Destination profiles are typically configured by the network administrator. At least one destination profile is required. You can configure multiple destination profiles of one or more types.

You can use one of the predefined destination profiles or define a desired profile. If you define a new profile, you must assign a profile name.



Note

If you use the Cisco AutoNotify service, the XML destination profile is required (see http://www.cisco.com/warp/customer/cc/serv/mkt/sup/tsssv/opmsup/smtan/anoti_ds.htm).

- Profile name—A string that uniquely identifies each user-defined destination profile and is limited to 32 alphanumeric characters. The format options for a user-defined destination profile are full-txt, short-txt, or XML (default).
- Destination address—The actual address, pertinent to the transport mechanism, to which the alert should be sent.
- Message formatting—The message format used for sending the alert (full text, short text, or XML).

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure predefined destination profile messaging options, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# callhome switch(config-callhome)# | Enters the Call Home submode. |
| Step 3 | switch(config-callhome)# destination-profile full-txt-destination email-addr person@place.com | Configures an e-mail address for the predefined full-txt-destination profile. The email addresses in this destination profile receives messages in full-txt format. The full-text format provides the complete, detailed explanation of the failure. Tip Use a standard e-mail address that does not have any text size restrictions. |
| | switch(config-callhome)# destination-profile full-txt-destination message-size 1000000 | Configures a maximum destination message size for the predefined full-txt-destination profile. The valid range is 0 to 1,000,000 bytes and the default is 500,000. A value of 0 implies that a message of any size can be sent. |
| Step 4 | switch(config-callhome)# destination-profile short-txt-destination email-addr person@place.com | Configures an e-mail address for the predefined short-txt-destination profile. The email-addresses in this destination profile receive messages in short-txt format. This format provides the basic explanation of the failure in the Call Home message. Tip Use a pager-related e-mail address for this option. |
| | switch(config-callhome)# destination-profile short-txt-destination message-size 100000 | Configures maximum destination message size for the predefined short-txt-destination profile. The valid range is 0 to 1,000,000 bytes and the default is 4000. A value of 0 implies that a message of any size can be sent. |
| Step 5 | switch(config-callhome)# destination-profile XML-destination email-addr findout@cisco.com | Configures an e-mail address for the predefined XML-destination profile. The email-addresses in this destination-profile receives messages in XML format. This format provides information that is compatible with Cisco Systems TAC support. Tip Do not add a pager-related e-mail address to this destination profile because of the large message size. |
| | switch(config-callhome)# destination-profile XML-destination message-size 100000 | Configures maximum destination message size for the predefined destination profile XML-destination. The valid range is 0 to 1,000,000 bytes and the default is 500,000. A value of 0 implies that a message of any size can be sent. |



Note

Steps 3, 4, and 5 in this procedure can be skipped or configured in any order.

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure a new destination-profile (and related parameters), follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# callhome switch(config-callhome)# | Enters the Call Home submode. |
| Step 3 | switch(config-callhome)# destination-profile test | Configures a new destination profile called test. |
| Step 4 | switch(config-callhome)# destination-profile test email-addr person@place.com | Configures the e-mail address for the user-defined destination profile (test) sent in default XML format. |
| Step 5 | switch(config-callhome)# destination-profile test message-size 1000000 | Configures a maximum message size for the destination email addresses in the user-defined destination profile (test) sent in default XML format. The valid range is 0 to 1,000,000 bytes and the default is 500,000. A value of 0 implies that a message of any size can be sent. |
| Step 6 | switch(config-callhome)# destination-profile test format full-txt | Configures message-format for the user-defined destination profile (test) to be full text format. |
| | switch(config-callhome)# destination-profile test format short-txt | Configures message-format for the user-defined destination profile (test) to be short text format. |


Note

Steps 4, 5, and 6 in this procedure can be skipped or configured in any order.

Alert Groups

An alert group is a predefined subset of Call Home alerts supported in all switches in the Cisco MDS 9000 Family. Different types of Call Home alerts are grouped into different alert groups depending on their type. You can associate one or more alert groups to each profile as required by your network.

The **alert-group** option allows you to select the set of Call Home alerts to be received by a destination profile (predefined or user-defined).

You can associate a destination profile with multiple alert groups.


Note

A Call Home alert is sent to e-mail destinations in a destination profile only if that Call Home alert belongs to one of the alert groups associated with that destination profile

To configure alert group for a destination profile, follow these steps:

| | Command | Purpose |
|--------|---|----------------------------|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# callhome switch(config-callhome)# | Enters Call Home submode. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|---|--|
| Step 3 | <code>switch(config-callhome)# destination-profile test1 alert-group test</code> | Optional. Configures user-defined destination profile (test1) to receive all user-generated Callhome test notifications. |
| | <code>switch(config-callhome)# destination-profile short-txt-destination alert-group test</code> | Optional. Configures predefined short-text destination profile to receive all user-generated Callhome test notifications. |
| Step 4 | <code>switch(config-callhome)# destination-profile test1 alert-group all</code> | Optional. Configures user-defined destination profile (test1) to receive Call Home notifications for all events |
| | <code>switch(config-callhome)# destination-profile short-txt-destination alert-group all</code> | Optional. Configures predefined short-text destination message profile to receive Call Home notifications for all (default) events |
| Step 5 | <code>switch(config-callhome)# destination-profile test1 alert-group Cisco-TAC</code> | Optional. Configures user-defined destination message profile (test1) to receive Call Home notifications for events that are meant only for Cisco TAC, or the Auto-notify service. |
| | <code>switch(config-callhome)# destination-profile xml-destination alert-group Cisco-TAC</code> | Optional. Configures predefined XML destination message profile to receive Call Home notifications for events that are meant only for Cisco TAC or the auto-notify service. |
| Step 6 | <code>switch(config-callhome)# destination-profile test1 alert-group environmental</code> | Optional. Configures user-defined destination message profile (test1) to receive Call Home notifications for power, fan, and temperature-related events. |
| | <code>switch(config-callhome)# destination-profile short-txt-destination alert-group environmental</code> | Optional. Configures predefined short-text destination message profile to receive Call Home notifications for power, fan, and temperature-related events. |
| Step 7 | <code>switch(config-callhome)# destination-profile test1 alert-group inventory</code> | Optional. Configures user-defined destination message profile (test1) to receive Call Home notifications for inventory status events. |
| | <code>switch(config-callhome)# destination-profile short-txt-destination alert-group inventory</code> | Optional. Configures predefined short-text destination message profile to receive Call Home notifications for inventory status events. |
| Step 8 | <code>switch(config-callhome)# destination-profile test1 alert-group linecard-hardware</code> | Optional. Configures user-defined destination message profile (test1) to receive Call Home notifications for module-related events. |
| | <code>switch(config-callhome)# destination-profile short-txt-destination alert-group linecard-hardware</code> | Optional. Configures predefined short-text destination message profile to receive Call Home notifications for module-related events. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|---------|---|--|
| Step 9 | <code>switch(config-callhome)# destination-profile test1 alert-group supervisor-hardware</code> | Optional. Configures user-defined destination message profile (test1) to receive Call Home notifications for supervisor-related events. |
| | <code>switch(config-callhome)# destination-profile short-txt-destination alert-group supervisor-hardware</code> | Optional. Configures predefined short-text destination message profile to receive Call Home notifications for supervisor-related events. |
| Step 10 | <code>switch(config-callhome)# destination-profile test1 alert-group system</code> | Optional. Configures user-defined destination message profile (test1) to receive Call Home notifications for software-related events. |
| | <code>switch(config-callhome)# destination-profile short-txt-destination alert-group system</code> | Optional. Configures predefined short-text destination message profile to receive Call Home notifications for software-related events. |

Call Home Message Levels

The **message-level** option allows you to filter messages based on their level of urgency. Each destination profile (predefined and user-defined) is associated with a Call Home message level threshold. Any message with a value lower than the urgency threshold is not sent. The urgency level ranges from 0 (lowest level of urgency) to 9 (highest level of urgency), and the default is 0 (all messages are sent).

To configure message level settings for destination profiles, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | <code>switch# config t</code> | Enters configuration mode. |
| Step 2 | <code>switch(config)# callhome</code> <code>switch(config-callhome)#</code> | Enters Call Home submode. |
| Step 3 | <code>switch(config-callhome)# destination-profile test message-level 5</code> | Optional. Configures the message level urgency as 5 and above for the user-defined profile (test1). |
| | <code>switch(config-callhome)# no destination-profile oldtest message-level 7</code> | Removes a previously configured urgency level and reverts it to the default of 0 (all messages are sent). |

Send documentation comments to mdsfeedback-doc@cisco.com.

Syslog-based Alerts

As of Cisco SAN-OS Release 2.0(1b), you can configure the switch to send certain syslog messages as Call Home messages. A new alert group, **syslog-group-port**, is added to select syslog messages for the port facility. The Call Home application maps the syslog severity level to corresponding callhome severity level (see [Table 30-4](#)).

Whenever a syslog message is generated, the Call Home application sends a Call Home Message depending on the mapping between the destination profile and the alert group mapping and based on the severity level of the generated syslog message. To receive a syslog-based Call Home alert, you must associate a destination profile with the syslog alert groups (currently there is only one syslog alert group—**syslog-group-port**) and configure the appropriate message level.

To configure the syslog-group-port, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# callhome switch(config-callhome)# | Enters Call Home submode. |
| Step 3 | switch(config-callhome)# destination-profile short-txt-destination alert-group syslog-group-port | Configures the predefined destination profile (short-txt-destination) to receive Callhome Notifications corresponding to syslog messages for the port facility. |
| Step 4 | switch(config-callhome)# destination-profile short-txt-destination message-level 5 | Optional. Configures the predefined destination-profile (short-txt-destination) to not receive any Callhome message for a syslog message which has lower priority than syslog priority level— alert . |

RMON-based Alerts

As of Cisco SAN-OS Release 2.0(1b), you can configure the switch to send Call Home notifications corresponding to RMON alert triggers. All RMON-based Call Home messages have message level set to NOTIFY (2). The alert-group **rmon** is defined for all RMON-based Call Home alerts. To receive an RMON-based CallHome alert, you must associate a destination profile with the **rmon** alert group.

To configure RMON alert groups, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# callhome switch(config-callhome)# | Enters Call Home submode. |
| Step 3 | switch(config-callhome)# destination-profile xml-destination alert-group rmon | Optional. Configures a destination message profile (rmon_group) to send Call Home notifications for configured RMON messages. |

Send documentation comments to mdsfeedback-doc@cisco.com.

E-Mail Options

You can configure the from, reply-to, and return-receipt e-mail addresses. While most e-mail address configurations are optional, you must configure the SMTP server address for the Call Home functionality to work.

Configuring General E-Mail Options

To configure general e-mail options, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# callhome switch(config-callhome)# | Enters Call Home submode. |
| Step 3 | switch(config-callhome)# transport email from user@company1.com | Optional. Configures the from e-mail address. |
| Step 4 | switch(config-callhome)# transport email reply-to person@place.com | Optional. Configures the reply-to e-mail address to which all responses should be sent. |

Configuring SMTP Server and Ports

To configure the SMTP server and port, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# callhome switch(config-callhome)# | Enters Call Home submode. |
| Step 3 | switch(config-callhome)# transport email smtp-server 192.168.1.1 | Configures the DNS or IP address of the SMTP server to reach the server. The port usage defaults to 25 if no port is specified. |
| | switch(config-callhome)# transport email smtp-server 192.168.1.1 port 30 | |
| | | Note The port number is optional and, if required, may be changed depending on the server location. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Periodic Inventory Notification

As of Cisco SAN-OS Release 2.0(1b), you can configure the switch to periodically send a message with an inventory of all the software services currently enabled and running on the switch along with hardware inventory information. The inventory is modified each time the switch is restarted nondisruptively.

By default, this feature is disabled in all switches in the Cisco MDS 9000 Family. When you enable this feature without configuring an interval value, the Call Home message is sent every 7 days. This value ranges from 1 to 30 days.

To enable periodic inventory notification in a Cisco MDS 9000 Family switch, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# callhome switch(config-callhome)# | Enters the Call Home submenu. |
| Step 3 | switch(config-callhome)# periodic-inventory notification | Enables the periodic inventory notification feature. By default, the Call Home message is sent every 7 days. |
| | switch(config-callhome)# no periodic-inventory notification | Disables the periodic inventory notification feature (default). |
| Step 4 | switch(config-callhome)# periodic-inventory notification interval 15 | Configures the periodic inventory notification message to be sent every 15 days. This value ranges from 1 to 30 days. |
| | switch(config-callhome)# no periodic-inventory notification interval 15 | Defaults to using the factory default of sending a Call Home message every 7 days. |

Duplicate Message Throttle

As of Cisco SAN-OS Release 2.0(1b), you can configure a throttling mechanism to limit the number of Call Home messages received for same event. If the same message is sent multiple times from the switch within a short period of time say, you may be swamped with a large number of duplicate messages.

By default, this feature is enabled in all switches in the Cisco MDS 9000 Family. When enabled, if the number of messages sent exceeds the maximum limit of 30 messages within the 2-hour time frame, then further messages for that alert type are discarded within that time frame. You cannot modify the time frame or the message counter limit.

If 2 hours have elapsed since the first such message was sent and a new message has to be sent, then the new message is sent and the time frame is reset to the time when the new message was sent and the count is reset to 1.

Send documentation comments to mdsfeedback-doc@cisco.com.

To enable message throttling in a Cisco MDS 9000 Family switch, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# callhome switch(config-callhome)# | Enters the Call Home submode. |
| Step 3 | switch(config-callhome)# no duplicate-message throttle | Disables the duplicate message throttling feature. |
| | switch(config-callhome)# duplicate-message throttle | Enables the duplicate message throttling feature (default). |

Call Home Enable Function

Once you have configured the contact information, you must enable the Call Home function.

The **enable** command is required for the Call Home function to start operating.

To enable the Call Home function, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# callhome switch(config-callhome)# | Enters Call Home submode. |
| Step 3 | switch(config-callhome)# enable callhome enabled successfully switch(config-callhome)# | Enables the Call Home function. |
| | switch(config-callhome)# disable switch(config-callhome)# | Disables the Call Home function. When you disable the Call Home function, all input events are ignored. |
| | | Note Even if Call Home is disabled, basic information for each Call Home event is sent. |

Call Home Configuration Distribution

As of Cisco SAN-OS Release 2.0(1b), you can enable fabric distribution for all Cisco MDS switches in the fabric. When you perform Call Home configurations, and distribution is enabled, that configuration is distributed to all the switches in the fabric.



Note

The Switch priority and the Syscontact name are not distributed.

You automatically acquire a fabric-wide lock when you issue the first configuration command after you enabled distribution in a switch. The Call Home application uses the effective and pending database model to store or commit the commands based on your configuration. When you commit the configuration changes, the effective database is overwritten by the configuration changes in the pending database and all the switches in the fabric receive the same configuration. After making the configuration changes, you can choose to discard the changes by aborting the changes instead of committing them. In either case, the lock is released. Refer to [Chapter 9, “Using the CFS Infrastructure”](#) for more information on the CFS application.

Send documentation comments to mdsfeedback-doc@cisco.com.

To enable Call Home fabric distribution, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# callhome switch(config-callhome)# | Enters Call Home submenu. |
| Step 3 | switch(config-callhome)# distribute | Enables Call Home configuration distribution to all switches in the fabric. Acquires a fabric lock and stores all future configuration changes in the pending database. |
| | switch(config-callhome)# no distribute | Disables (default) Call Home configuration distribution to all switches in the fabric. |

To commit the Call Home configuration changes, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# callhome switch(config-callhome)# | Enters Call Home submenu. |
| Step 3 | switch(config-callhome)# commit | Distributes the configuration changes to all switches in the fabric and releases the lock. Overwrites the effective database with the changes made to the pending database. |

To discard the Call Home configuration changes, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# callhome switch(config-callhome)# | Enters Call Home submenu. |
| Step 3 | switch(config-callhome)# abort | Discards the configuration changes in the pending database and releases the fabric lock. |

Fabric Lock Override

If you have performed a Call Home task and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.



Tip

The changes are only available in the volatile directory and are subject to being discarded if the switch is restarted.

To use administrative privileges and release a locked Call Home session, use the **clear callhome session** command.

```
switch# clear callhome session
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Database Merge Guidelines

Refer to the “[CFS Merge Support](#)” section on page 9-7 for detailed concepts.

When merging two Call Home databases, follow these guidelines:

- Be aware that the merged database contains the following information:
 - A superset of all the destination profiles from the dominant and subordinate switches take part in the merge protocol.
 - The e-mail addresses and alert groups for the destination profiles.
 - Other configuration information (for example, message throttling, periodic inventory) from the switch which existed in the dominant switch before the merge.
- Verify that two destination profiles do not have the same name (even if they have different configuration information) on the subordinate and dominant switches. If they do contain the same name, the merge operation will fail. You must then modify or delete the conflicting destination profile on the required switch.

Call Home Communications Test

Use the **test** command to simulate a message generation.

To test the Call Home function, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# callhome test trying to send test callhome message successfully sent test callhome message | Sends a test message to the configured destination(s). |
| Step 2 | switch# callhome test inventory trying to send test callhome message successfully sent test callhome message | Sends a test inventory message to the configured destination(s). |

Displaying Call Home Information

Use the **show callhome** command to display the configured Call Home information (see Examples 30-1 to 30-7).

Example 30-1 Displays Configured Call Home Information

```
switch# show callhome
callhome enabled
Callhome Information:
contact person name:who@where
contact person's email:person@place.com
contact person's phone number:310-408-4000
street addr:1234 Picaboo Street, Any city, Any state, 12345
site id:Site1ManhattanNewYork
customer id:Customer1234
contract id:Cisco1234
switch priority:0
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 30-2 Displays Information for All Destination Profiles (Predefined and User-Defined)

```
switch# show callhome destination-profile
XML destination profile information
maximum message size:500000
message format:XML
message-level:0
email addresses configured:
alert groups configured:
cisco_tac

test destination profile information
maximum message size:100000
message format:full-txt
message-level:5
email addresses configured:
cchetty@isco.com

alert groups configured:
test

full-txt destination profile information
maximum message size:500000
message format:full-txt
message-level:0
email addresses configured:

alert groups configured:
all

short-txt destination profile information
maximum message size:4000
message format:short-txt
message-level:0
email addresses configured:

alert groups configured:
all
```

Example 30-3 Displays Information for a User-defined Destination Profile

```
switch# show callhome destination-profile test
test destination profile information
maximum message size:100000
message format:full-txt
message-level:5
email addresses configured:
user@company.com

alert groups configured:
test
```

Example 30-4 Displays the Full-Text Profile

```
switch# show callhome destination-profile profile full-txt-destination
full-txt destination profile information
maximum message size:250000
email addresses configured:
person2@company2.com
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 30-5 Displays the Short-Text Profile

```
switch# show callhome destination-profile profile short-txt-destination
Short-txt destination profile information
maximum message size:4000
email addresses configured:
person2@company2.com
```

Example 30-6 Displays the XML Destination Profile

```
switch# show callhome destination-profile profile XML-destination
XML destination profile information
maximum message size:250000
email addresses configured:
findout@cisco.com
```

Example 30-7 Displays E-Mail and SMTP Information

```
switch# show callhome transport-email
from email addr:user@company1.com
reply to email addr:pointer@company.com
return receipt email addr:user@company1.com
smtp server:server.company.com
smtp server port:25
```

Sample Syslog Alert Notification in Full-txt Format

```
source:MDS9000
Switch Priority:7
Device Id:DS-C9506@C@FG@07120011
Customer Id:basu
Contract Id:123
Site Id:Bangalore
Server Id:DS-C9506@C@FG@07120011
Time of Event:2004-10-08T11:10:44
Message Name:SYSLOG_ALERT
Message Type:Syslog
Severity Level:2
System Name:10.76.100.177
Contact Name:Basavaraj B
Contact Email:bbendige@cisco.com
Contact Phone:+91-80-310-1718
Street Address:#71 , Miller's Road
Event Description:2004 Oct 8 11:10:44 10.76.100.177 %PORT-5-IF_TRUNK_UP: %$VSAN 1%$
Interface fc2/5, vsan 1 is up

syslog_facility:PORT
start chassis information:
Affected Chassis:DS-C9506
Affected Chassis Serial Number:FG@07120011
Affected Chassis Hardware Version:0.104
Affected Chassis Software Version:2.0(1)
Affected Chassis Part No:73-8607-01
end chassis information:
```


Send documentation comments to mdsfeedback-doc@cisco.com.

Sample Syslog Alert Notification in XML Format

```
X-Mozilla-Status2: 02000000
Return-Path: <tester@cisco.com>
...

<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<!DOCTYPE mml SYSTEM "mml10.dtd">
<!--
Alert:SYSLOG_ALERT
-->
<mml>
<header>
<time>2004-09-30T06:12:36</time>
<name>SYSLOG_ALERT</name>
<type>Syslog</type>
<level>2</level>
<source>MDS9000</source>
<priority>7</priority>
<deviceId>DS-C9506@C@FOX0712S00H</deviceId>
<custId>911</custId>
<contractId>33445</contractId>
<siteId>91111</siteId>
<serverId>DS-C9506@C@FOX0712S00H</serverId>
</header>
<body>
<msgDesc>2004 Sep 30 06:12:36 switch186 %PORT-5-IF_UP: %$VSAN 2000%$ Interface fc1/10 is
up in mode FL
</msgDesc>
<sysName>switch186</sysName>
<sysContact>USA</sysContact>
<sysContactEmail>billgates@microsoft.com</sysContactEmail>
<sysContactPhoneNumber>+91-080-8888888</sysContactPhoneNumber>
<sysStreetAddress>91</sysStreetAddress>
<chassis>
<name>DS-C9506</name>
<serialNo>FOX0712S00H</serialNo>
<partNo>73-8697-01</partNo>
<hwVersion>0.104</hwVersion>
<swVersion>2.0(1)</swVersion>
</chassis>
<nvp>
<name>syslog_facility</name>
<value>PORT</value>
</nvp>
</body>
</mml>
```

Sample RMON Notification in XML Format

```
Return-Path: <tester@cisco.com>
...
<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<!DOCTYPE mml SYSTEM "mml10.dtd">
<!--
Alert:RMON_ALERT
-->
<mml>
<header>
<time>2004-10-12T04:59:13</time>
<name>RMON_ALERT</name>
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
<type>RMON</type>
<level>2</level>
<source>MDS9000</source>
<priority>3</priority>
<deviceId>DS-C9506@C@FOX0712S00H</deviceId>
<custId>0</custId>
<contractId>u</contractId>
<siteId>&lt;/siteId>
<serverId>DS-C9506@C@FOX0712S00H</serverId>
</header>
<body>
<msgDesc>rlaxmina-w2k07</msgDesc>
<sysName>switch186</sysName>
<sysContact>USA</sysContact>
<sysContactEmail>billgates@microsoft.com</sysContactEmail>
<sysContactPhoneNumber>+91-080-000000</sysContactPhoneNumber>
<sysStreetAddress>91</sysStreetAddress>
<chassis>
<name>DS-C9506</name>
<serialNo>FOX0712S00H</serialNo>
<partNo>73-8697-01</partNo>
<hwVersion>0.104</hwVersion>
<swVersion>2.0(1)</swVersion>
</chassis>
<nvp>
<name>ThresholdType</name>
<value>RisingThreshold</value>
</nvp>
<nvp>
<name>ThresholdValue</name>
<value>0</value>
</nvp>
<nvp>
<name>AlarmValue</name>
<value>0</value>
</nvp>
</body>
</mml>
```

Default Settings

Table 30-1 lists the default Call Home default settings.

Table 30-1 **Default Call Home Settings**

| Parameters | Default |
|---|-----------|
| Destination message size for a message sent in full text format. | 500,000. |
| Destination message size for a message sent in XML format. | 500,000. |
| Destination message size for a message sent in short text format. | 4,000. |
| DNS or IP address of the SMTP server to reach the server if no port is specified. | 25. |
| Alert group association with profile. | All. |
| Format type. | XML. |
| Call Home message level. | 0 (zero). |

Send documentation comments to mdsfeedback-doc@cisco.com.

Event Triggers

This section discusses Call Home trigger events. Trigger events are divided into categories, with each category assigned commands to execute when the event occurs. The command output is included in the transmitted message. [Table 30-2](#) lists the trigger events.

Table 30-2 **Event Triggers**

| Event | Alert Group | Event Name | Description | Call Home Message Level |
|-----------|--|------------------------------|--|-------------------------|
| Call Home | System and CISCO_TAC | SW_CRASH | A software process has crashed with a stateless restart, indicating an interruption of a service. | 5 |
| | System and CISCO_TAC | SW_SYSTEM_INCONSISTENT | Inconsistency detected in software or file system. | 5 |
| | Environmental and CISCO_TAC | TEMPERATURE_ALARM | Thermal sensor indicates temperature reached operating threshold. | 6 |
| | | POWER_SUPPLY_FAILURE | Power supply failed. | 6 |
| | | FAN_FAILURE | Cooling fan has failed. | 5 |
| | Switching module and CISCO_TAC | LINECARD_FAILURE | Switching module operation failed. | 7 |
| | | POWER_UP_DIAGNOSTICS_FAILURE | Switching module failed power-up diagnostics. | 7 |
| | Line Card Hardware and CISCO_TAC | PORT_FAILURE | Hardware failure of interface port(s). | 6 |
| | Line Card Hardware, Supervisor Hardware, and CISCO_TAC | BOOTFLASH_FAILURE | Failure of boot compact Flash card. | 6 |
| | Supervisor module and CISCO_TAC | SUP_FAILURE | Supervisor module operation failed. | 7 |
| | | POWER_UP_DIAGNOSTICS_FAILURE | Supervisor module failed power-up diagnostics. | 7 |
| | Supervisor Hardware and CISCO_TAC | INBAND_FAILURE | Failure of in-band communications path. | 7 |
| | Supervisor Hardware and CISCO_TAC | EOBC_FAILURE | Ethernet out-of-band channel communications failure. | 6 |
| | Supervisor Hardware and CISCO_TAC | MGMT_PORT_FAILURE | Hardware failure of management Ethernet port. | 5 |
| | License | LICENSE_VIOLATION | Feature in use is not licensed (Cisco MDS SAN-OS Release 1.3), and are turned off after grace period expiration. | 6 |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 30-2 **Event Triggers (continued)**

| Event | Alert Group | Event Name | Description | Call Home Message Level |
|-------------|-------------------------|--------------------|---|-------------------------|
| Inventory | Inventory and CISCO_TAC | COLD_BOOT | Switch is powered up and reset to a cold boot sequence. | 2 |
| | | HARDWARE_INSERTION | New piece of hardware inserted into the chassis. | 2 |
| | | HARDWARE_REMOVAL | Hardware removed from the chassis. | 2 |
| Test | Test and CISCO_TAC | TEST | User generated test. | 2 |
| Port syslog | Syslog-group-port | SYSLOG_ALERT | Syslog messages corresponding to the port facility. | 2 |
| RMON | RMON | RMON_ALERT | RMON alert trigger messages. | 2 |

Table 30-3 lists event categories and command outputs.

Table 30-3 **Event Categories and Executed Commands**

| Event Category | Description | Executed Commands |
|---------------------------|--|--|
| System | Events generated by failure of a software system that is critical to unit operation. | show tech-support show system redundancy status |
| Environmental | Events related to power, fan, and environment sensing elements such as temperature alarms. | show module show environment |
| Switching module hardware | Events related to standard or intelligent switching modules. | show tech-support |
| Supervisor hardware | Events related to supervisor modules. | show tech-support |
| Inventory | Inventory status is provided whenever a unit is cold booted, or when FRUs are inserted or removed. This is considered a noncritical event, and the information is used for status and entitlement. | show version |
| Test | User generated test message. | show version |

Call Home Message Levels

Call Home messages (sent for syslog alert groups) have the syslog severity level mapped to the Call Home message level.

This section discusses the severity levels for a Call Home message when using one or more switches in the Cisco MDS 9000 Family. Call Home message levels are preassigned per event type.



Note

Call Home severity levels are not the same as system message logging severity levels (see [Chapter 36](#), “Configuring System Message Logging”).

Severity levels range from 0 to 9, with 9 having the highest urgency. Each syslog level has keywords and a corresponding syslog level as listed in [Table 30-4](#).

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 30-4 Severity and Syslog Level Mapping

| Call Home Level | Keyword Used | Syslog Level | Description |
|------------------|---------------------|-----------------|--|
| Catastrophic (9) | Catastrophic | N/A | Network wide catastrophic failure. |
| Disaster (8) | Disaster | N/A | Significant network impact. |
| Fatal (7) | Fatal | Emergency (0) | System is unusable. |
| Critical (6) | Critical | Alert (1) | Critical conditions, immediate attention needed. |
| Major (5) | Major | Critical (2) | Major conditions. |
| Minor (4) | Minor | Error (3) | Minor conditions. |
| Warning (3) | Warning | Warning (4) | Warning conditions. |
| Notify (2) | Notification | Notice (5) | Basic notification and informational messages. Possibly independently insignificant. |
| Normal (1) | Normal | Information (6) | Normal event signifying return to normal state. |
| Debug (0) | Debugging | Debug (7) | Debugging messages. |

Message Contents

The following contact information can be configured on the switch:

- Name of the contact person
- Phone number of the contact person
- E-mail address of the contact person
- Mailing address to which replacement parts must be shipped, if required
- Site ID of the network where the site is deployed
- Contract ID to identify the service contract of the customer with the service provider

[Table 30-5](#) describes the short text formatting option for all message types.

Table 30-5 Short Text Messages

| Data Item | Description |
|-------------------------|--|
| Device identification | Configured device name |
| Date/time stamp | Time stamp of the triggering event |
| Error isolation message | Plain English description of triggering event |
| Alarm urgency level | Error level such as that applied to system message |

[Table 30-6](#), [Table 30-7](#), and [Table 30-8](#) display the information contained in plain text and XML messages.

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 30-6 *Reactive Event Message Format*

| Data Item (Plain text and XML) | Description (Plain text and XML) | XML Tag (XML only) |
|-----------------------------------|---|---------------------------|
| Time stamp | Date and time stamp of event in ISO time notation: YYYY-MM-DDTHH:MM:SS. Note The time zone or daylight savings time (DST) offset from UTC has already been added or subtracted. T is the hardcoded limiter for the time. | /mml/header/time |
| Message name | Name of message. Specific event names are listed in the “ Event Triggers ” section on page 30-19. | /mml/header/name |
| Message type | Specifically “Call Home.” | /mml/header/type |
| Message group | Specifically “reactive.” | /mml/header/group |
| Severity level | Severity level of message (see Table 30-4). | /mml/header/level |
| Source ID | Product type for routing. | /mml/header/source |
| Device ID | Unique device identifier (UDI) for end device generating message. This field should empty if the message is non-specific to a fabric switch. Format: type@Sid@serial, where <ul style="list-style-type: none"> Type is the product model number from backplane SEEPROM. @ is a separator character. Sid is “C” identifying serial ID as a chassis serial number. Serial number as identified by the Sid field. Example: “DS-C9509@C@12345678” | /mml/ header/deviceId |
| Customer ID | Optional user-configurable field used for contract info or other ID by any support service. | /mml/ header/customerID |
| Contract ID | Optional user-configurable field used for contract info or other ID by any support service. | /mml/ header /contractId |
| Site ID | Optional user-configurable field used for Cisco-supplied site ID or other data meaningful to alternate support service. | /mml/ header/siteId |
| Server ID | If the message is generated from the fabric switch, it is the unique device identifier (UDI) of the switch. Format: type@Sid@serial, where <ul style="list-style-type: none"> Type is the product model number from backplane SEEPROM. @ is a separator character. Sid is “C” identifying serial ID as a chassis serial number. Serial number as identified by the Sid field. Example: “DS-C9509@C@12345678” | /mml/header/serverId |
| Message description | Short text describing the error. | /mml/body/msgDesc |
| Device name | Node that experienced the event. This is the host name of the device. | /mml/body/sysName |
| Contact name | Name of person to contact for issues associated with the node experiencing the event. | /mml/body/sysContact |
| Contact e-mail | E-mail address of person identified as contact for this unit. | /mml/body/sysContactEmail |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 30-6 *Reactive Event Message Format (continued)*

| Data Item (Plain text and XML) | Description (Plain text and XML) | XML Tag (XML only) |
|---|--|--|
| Contact phone number | Phone number of the person identified as the contact for this unit. | /mml/body/sysContactPhone Number |
| Street address | Optional field containing street address for RMA part shipments associated with this unit. | /mml/body/sysStreetAddress |
| Model name | Model name of the switch. This is the specific model as part of a product family name. | /mml/body/chassis/name |
| Serial number | Chassis serial number of the unit. | /mml/body/chassis/serialNo |
| Chassis part number | Top assembly number of the chassis. | /mml/body/chassis/partNo |
| Chassis hardware version | Hardware version of chassis. | /mml/body/chassis/hwVersion |
| Supervisor module software version | Top level software version. | /mml/body/chassis/swVersion |
| Affected FRU name | Name of the affected FRU generating the event message. | /mml/body/fru/name |
| Affected FRU serial number | Serial number of affected FRU. | /mml/body/fru/serialNo |
| Affected FRU part number | Part number of affected FRU. | /mml/body/fru/partNo |
| FRU slot | Slot number of FRU generating the event message. | /mml/body/fru/slot |
| FRU hardware version | Hardware version of affected FRU. | /mml/body/fru/hwVersion |
| FRU software version | Software version(s) running on affected FRU. | /mml/body/fru/swVersion |
| Command output name | The exact name of the issued command. | /mml/attachments/attachment/ name |
| Attachment type | Specifically command output. | /mml/attachments/attachment/ type |
| MIME type | Normally text or plain or encoding type. | /mml/attachments/attachment/ mime |
| Command output text | Output of command automatically executed (see Table 30-3). | /mml/attachments/attachment/ atdata |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 30-7 Inventory Event Message Format

| Data Item (Plain text and XML) | Description (Plain text and XML) | XML Tag (XML only) |
|-----------------------------------|---|---------------------------|
| Time stamp | Date and time stamp of event in ISO time notation: YYYY-MM-DDTHH:MM:SS. Note The time zone or daylight savings time (DST) offset from UTC has already been added or subtracted. T is the hardcoded limiter for the time. | /mml/header/time |
| Message name | Name of message. Specifically “Inventory Update” Specific event names are listed in the “Event Triggers” section on page 30-19 . | /mml/header/name |
| Message type | Specifically “Inventory Update”. | /mml/header/type |
| Message group | Specifically “proactive”. | /mml/header/group |
| Severity level | Severity level of inventory event is level 2 (see Table 30-4). | /mml/header/level |
| Source ID | Product type for routing at Cisco. Specifically “MDS 9000” | /mml/header/source |
| Device ID | Unique Device Identifier (UDI) for end device generating message. This field should empty if the message is non-specific to a fabric switch. Format: type@Sid@serial, where <ul style="list-style-type: none"> Type is the product model number from backplane SEEPROM. @ is a separator character. Sid is “C” identifying serial ID as a chassis serial number. Serial: The serial number as identified by the Sid field. Example: “DS-C9509@C@12345678” | /mml/ header /deviceId |
| Customer ID | Optional user-configurable field used for contact info or other ID by any support service. | /mml/ header /customerID |
| Contract ID | Optional user-configurable field used for contact info or other ID by any support service. | /mml/ header /contractId |
| Site ID | Optional user-configurable field, can be used for Cisco-supplied site ID or other data meaningful to alternate support service. | /mml/ header /siteId |
| Server ID | If the message is generated from the fabric switch, it is the Unique device identifier (UDI) of the switch. Format: type@Sid@serial, where <ul style="list-style-type: none"> Type is the product model number from backplane SEEPROM. @ is a separator character. Sid is “C” identifying serial ID as a chassis serial number. Serial: The serial number as identified by the Sid field. Example: “DS-C9509@C@12345678” | /mml/header/serverId |
| Message description | Short text describing the error. | /mml/body/msgDesc |
| Device name | Node that experienced the event. | /mml/body/sysName |
| Contact name | Name of person to contact for issues associated with the node experiencing the event. | /mml/body/sysContact |
| Contact e-mail | E-mail address of person identified as contact for this unit. | /mml/body/sysContactEmail |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 30-7 *Inventory Event Message Format (continued)*

| Data Item (Plain text and XML) | Description (Plain text and XML) | XML Tag (XML only) |
|---|---|------------------------------------|
| Contact phone number | Phone number of the person identified as the contact for this unit. | /mml/body/sysContactPhoneNumber |
| Street address | Optional field containing street address for RMA part shipments associated with this unit. | /mml/body/sysStreetAddress |
| Model name | Model name of the unit. This is the specific model as part of a product family name. | /mml/body/chassis/name |
| Serial number | Chassis serial number of the unit. | /mml/body/chassis/serialNo |
| Chassis part number | Top assembly number of the chassis. | /mml/body/chassis/partNo |
| Chassis hardware version | Hardware version of chassis. | /mml/body/chassis/hwVersion |
| Supervisor module software version | Top level software version. | /mml/body/chassis/swVersion |
| FRU name | Name of the affected FRU generating the event message. | /mml/body/fru/name |
| FRU s/n | Serial number of FRU. | /mml/body/fru/serialNo |
| FRU part number | Part number of FRU. | /mml/body/fru/partNo |
| FRU slot | Slot number of FRU. | /mml/body/fru/slot |
| FRU hardware version | Hardware version of FRU. | /mml/body/fru/hwVersion |
| FRU software version | Software version(s) running on FRU. | /mml/body/fru/swVersion |
| Command output name | The exact name of the issued command. | /mml/attachments/attachment/name |
| Attachment type | Specifically command output. | /mml/attachments/attachment/type |
| MIME type | Normally text or plain or encoding type. | /mml/attachments/attachment/mime |
| Command output text | Output of command automatically executed after event categories (see “Event Triggers” section on page 30-19). | /mml/attachments/attachment/atdata |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 30-8 **User-Generated Test Message Format**

| Data Item (Plain text and XML) | Description (Plain text and XML) | XML Tag (XML only) |
|-----------------------------------|---|--------------------------|
| Time stamp | Date and time stamp of event in ISO time notation: YYYY-MM-DDTHH:MM:SS. Note The time zone or daylight savings time (DST) offset from UTC has already been added or subtracted. T is the hardcoded limiter for the time. | /mml/header/time |
| Message name | Name of message. Specifically test message for test type message. Specific event names listed in the “Event Triggers” section on page 30-19 . | /mml/header/name |
| Message type | Specifically “Test Call Home”. | /mml/header/type |
| Message group | This field should be ignored by the receiving Call Home processing application, but may be populated with either “proactive” or “reactive”. | /mml/header/group |
| Severity level | Severity level of message, test Call Home message (see Table 30-4). | /mml/header/level |
| Source ID | Product type for routing. | /mml/header/source |
| Device ID | Unique device identifier (UDI) for end device generating message. This field should empty if the message is non-specific to a fabric switch. Format: type@Sid@serial, where <ul style="list-style-type: none"> Type is the product model number from backplane SEEPROM. @ is a separator character. Sid is “C” identifying serial ID as a chassis serial number. Serial: The serial number as identified by the Sid field. Example: “DS-C9509@C@12345678” | /mml/ header /deviceId |
| Customer ID | Optional user-configurable field used for contract info or other ID by any support service. | /mml/ header /customerId |
| Contract ID | Optional user-configurable field used for contract info or other ID by any support service. | /mml/ header /contractId |
| Site ID | Optional user-configurable field used for Cisco-supplied site ID or other data meaningful to alternate support service. | /mml/ header /siteId |
| Server ID | If the message is generated from the fabric switch, it is the Unique device identifier (UDI) of the switch. Format: type@Sid@serial, where <ul style="list-style-type: none"> Type is the product model number from backplane SEEPROM. @ is a separator character. Sid is “C” identifying serial ID as a chassis serial number. Serial: The serial number as identified by the Sid field. Example: “DS-C9509@C@12345678” | /mml/header/serverId |
| Message description | Short text describing the error. | /mml/body/msgDesc |
| Device name | Switch that experienced the event. | /mml/body/sysName |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 30-8 ***User-Generated Test Message Format (continued)***

| Data Item (Plain text and XML) | Description (Plain text and XML) | XML Tag (XML only) |
|---|--|--|
| Contact name | Name of person to contact for issues associated with the node experiencing the event. | /mml/body/sysContact |
| Contact Email | E-mail address of person identified as contact for this unit. | /mml/body/sysContactEmail |
| Contact phone number | Phone number of the person identified as the contact for this unit. | /mml/body/sysContactPhone Number |
| Street address | Optional field containing street address for RMA part shipments associated with this unit. | /mml/body/sysStreetAddress |
| Model name | Model name of the switch. This is the specific model as part of a product family name. | /mml/body/chassis/name |
| Serial number | Chassis serial number of the unit. | /mml/body/chassis/serialNo |
| Chassis part number | Top assembly number of the chassis. For example, 800-xxx-xxxx. | /mml/body/chassis/partNo |
| Command output text | Output of command automatically executed after event categories listed in Table 30-3 . | /mml/attachments/attachmen t/atdata |
| MIME type | Normally text or plain or encoding type. | /mml/attachments/attachmen t/mime |
| Attachment type | Specifically command output. | /mml/attachments/attachmen t/type |
| Command output name | The exact name of the issued command. | /mml/attachments/attachmen t/name |

Send documentation comments to mdsfeedback-doc@cisco.com.



Configuring Domain Parameters

The Fibre Channel domain (fcdomain) feature performs principal switch selection, domain ID distribution, FC ID allocation, and fabric reconfiguration functions as described in the FC-SW-2 standards. The domains are configured on a per VSAN basis. If you do not configure a domain ID, the local switch uses a random ID.



Caution

Changes to fcdomain parameters should not be performed on a daily basis. These changes should be made by an administrator or individual who is completely familiar with switch operations.



Tip

When you change the configuration, be sure to save the running configuration. The next time you reboot the switch, the saved configuration is used. If you do not save the configuration, the previously saved startup configuration is used.

This chapter includes the following sections:

- [About fcdomain Phases, page 31-2](#)
- [Domain Restart, page 31-3](#)
- [Domain Configuration, page 31-4](#)
- [Switch Priority, page 31-6](#)
- [Allowed Domain ID Lists, page 31-6](#)
- [Merged Stable Fabrics, page 31-7](#)
- [Contiguous Domain Assignments, page 31-7](#)
- [fcdomain Initiation, page 31-8](#)
- [Fabric Name, page 31-8](#)
- [Incoming RCFs, page 31-9](#)
- [Persistent FC IDs, page 31-9](#)
- [Persistent FC IDs Manual Configuration, page 31-10](#)
- [Persistent FC ID Selective Purging, page 31-13](#)
- [Displaying fcdomain Information, page 31-13](#)
- [Default Settings, page 31-17](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

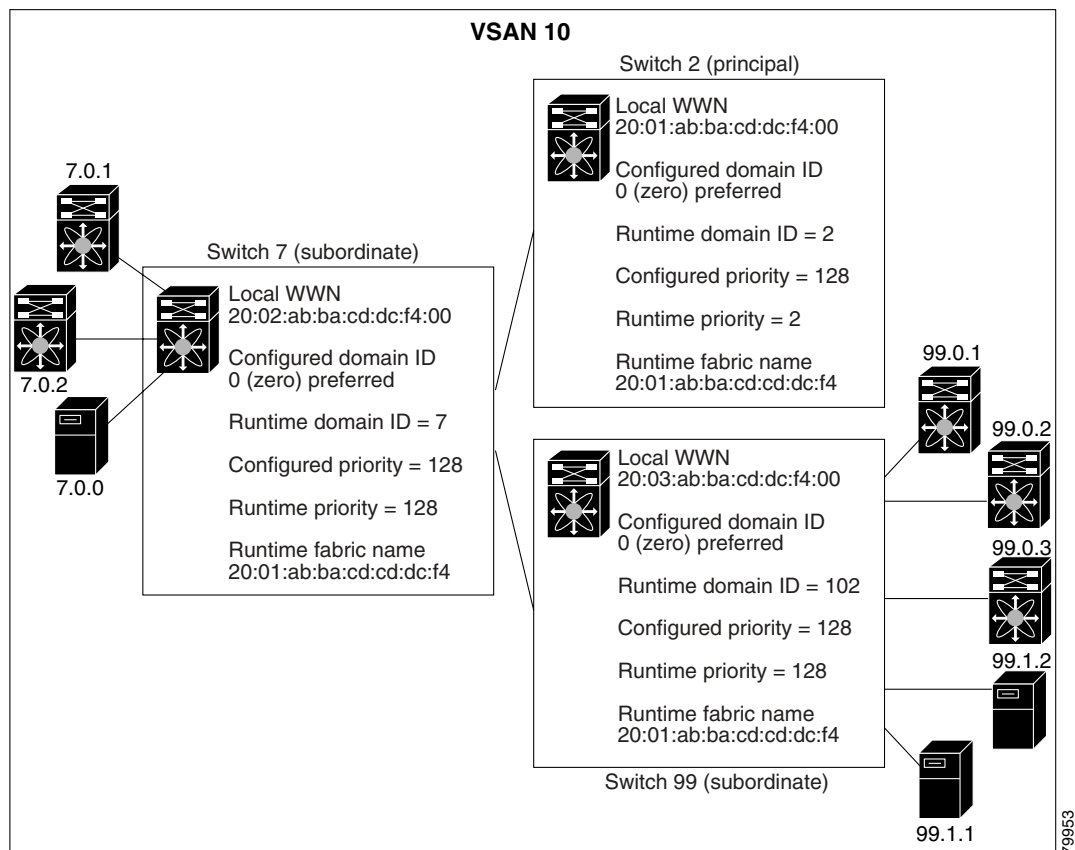
About fcdomain Phases

This section describes each fcdomain phase:

- Principal switch selection—This phase guarantees the selection of a unique principal switch across the fabric.
- Domain ID distribution—This phase guarantees each switch in the fabric obtains a unique domain ID.
- FC ID allocation—This phase guarantees a unique FC ID assignment to each device attached to the corresponding switch in the fabric.
- Fabric reconfiguration—This phase guarantees a resynchronization of all switches in the fabric to ensure they simultaneously restart a new principal switch selection phase.

See [Figure 31-1](#).

Figure 31-1 Sample fcdomain Configuration



Note

Domain IDs and VSAN values used in all procedures are only provided as examples. Be sure to use IDs and values that apply to your configuration.

Send documentation comments to mdsfeedback-doc@cisco.com.

Domain Restart

Fibre Channel domains can be started disruptively or nondisruptively. If you perform a disruptive restart, reconfigure fabric (RCF) frames are sent to other switches in the fabric. If you perform a nondisruptive restart, build fabric (BF) frames are sent to other switches in the fabric.



Note

A static domain is specifically configured by the user and may be different from the runtime domain. If the domain IDs are different, the runtime domain ID changes to take on the static domain ID after the next restart.



Tip

If a VSAN is in interop mode, you cannot restart the fcdomain for that VSAN disruptively.

You can apply most of the configurations to their corresponding runtime values. Each of the following sections provide further details on how the fcdomain parameters are applied to the runtime values.

The **fcdomain restart** command applies your changes to the runtime settings. Use the **restart disruptive** option to apply most of the configurations to their corresponding runtime values.

To restart the fabric disruptively or nondisruptively, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fcdomain restart vsan 1 | Forces the VSAN to reconfigure without traffic disruption. |
| | switch(config)# fcdomain restart disruptive vsan 1 | Forces the VSAN to reconfigure with data traffic disruption. |

Send documentation comments to mdsfeedback-doc@cisco.com.

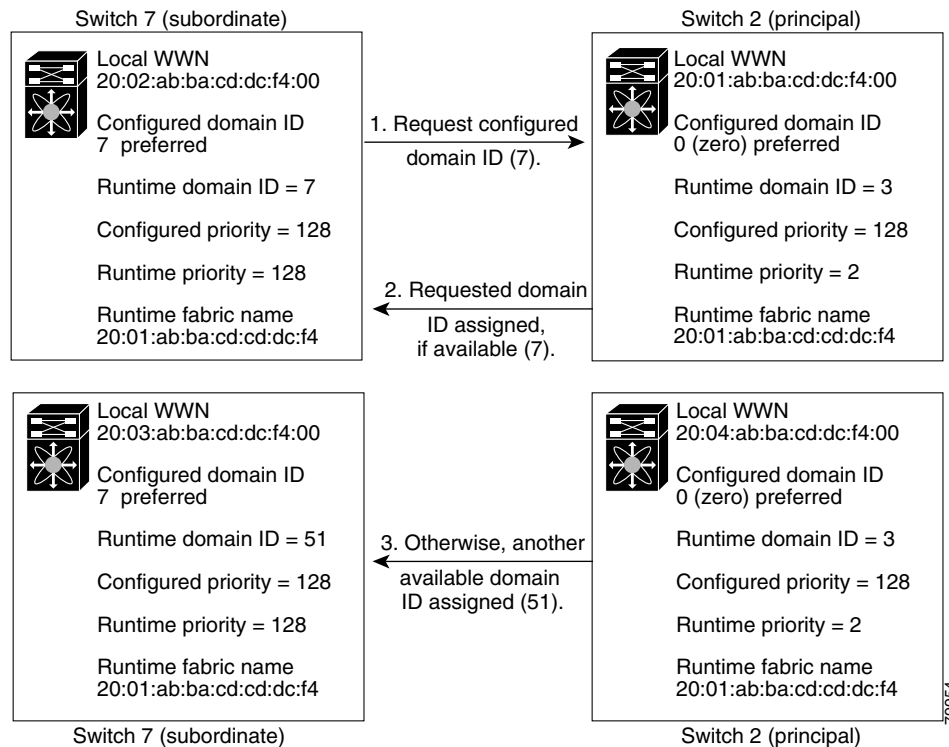
Domain Configuration

The configured domain ID can be **preferred** or **static**. By default, the configured domain is **0** and the configured type is **preferred**. If you do not configure a domain ID, the local switch sends a random ID in its request.

When a subordinate switch requests a domain, the following process takes place (see Figure 31-2):

1. The local switch sends a configured domain ID request to the principal switch.
2. The principal switch assigns the requested domain ID if available. Otherwise, it assigns another available domain ID.

Figure 31-2 Configuration Process Using the preferred Option



The behavior for a subordinate switch changes based on the allowed domain ID lists, on the configured domain ID, and on the domain ID that the principal switch has assigned to the requesting switch.

- When the received domain ID is not within the allowed list, the requested domain ID becomes the runtime domain ID and all interfaces on that VSAN are isolated.
- When the assigned and requested domain IDs are the same, the **preferred** and **static** options are not relevant, and the assigned domain ID becomes the runtime domain ID.
- When the assigned and requested domain IDs are different, the following cases apply:
 - If the configured type is **static**, the assigned domain ID is discarded, all local interfaces are isolated, and the local switch assigns itself the configured domain ID, which becomes the runtime domain ID.
 - If the configured type is **preferred**, the local switch accepts the domain ID assigned by the principal switch and the assigned domain ID becomes the runtime domain ID.

Send documentation comments to mdsfeedback-doc@cisco.com.

If you change the configured domain ID, the change is only accepted if the new domain ID is included in all the allowed domain ID lists currently configured in the VSAN. Alternatively, you can also configure zero-preferred domain ID.



Note

The 0 (zero) value can be configured only if you use the **preferred** option.

While the **static** option can be applied to runtime after a disruptive or nondisruptive restart, the **preferred** option is applied to runtime only after a disruptive restart (see the “[Domain Restart](#)” section on page 31-3).



Tip

When the FICON feature is enabled in a given VSAN, the domain ID for that VSAN remains in the static state. You can change the static ID value but you cannot change it to the preferred option.



Caution

You must issue the **fcdomain restart** command if you want to apply the configured domain changes to the runtime domain.

To specify a **preferred** or a **static** domain ID, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fcdomain domain 3 preferred vsan 8 | Configures the switch in VSAN 8 to request a preferred domain ID 3 and accepts any value assigned by the principal switch. |
| | switch(config)# no fcdomain domain 3 preferred vsan 8 | Resets the configured domain ID to 0 (default) in VSAN 8. The configured domain ID becomes 0 preferred. |
| Step 3 | switch(config)# fcdomain domain 2 static vsan 237 | Configures the switch in VSAN 237 to accept only a specific value and moves the local interfaces in VSAN 237 to an isolated state if the requested domain ID is not granted. |
| | switch(config)# no fcdomain domain 18 static vsan 237 | Resets the configured domain ID to factory defaults in VSAN 237. The configured domain ID becomes 0 preferred. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Switch Priority

By default, the configured priority is 128. The valid range to set the priority is between 1 and 254. Priority 1 has the highest priority. Value 255 is accepted from other switches, but cannot be locally configured.

Any new switch cannot become the principal switch when it joins a stable fabric. During the principal switch selection phase, the switch with the highest priority becomes the principal switch. If two switches have the same configured priority, the switch with the lower WWN becomes the principal switch.

The priority configuration is applied to runtime when the fcdomain is restarted (see the “[Domain Restart](#)” section on page 31-3). This configuration is applicable to both disruptive and nondisruptive restarts.

To configure the priority for the principal switch, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fcdomain priority 25 VSAN 99 | Configures a priority of 25 for the local switch in VSAN 99. |
| | switch(config)# no fcdomain priority 25 VSAN 99 | Reverts the priority to the factory default (128) in VSAN 99. |

Allowed Domain ID Lists

By default, the valid range for an assigned domain ID list is from 1 to 239. You can specify a list of ranges to be in the allowed domain ID list and separate each range with a comma. The principal switch assigns domain IDs that are available in the locally-configured allowed domain list.



Tip

If you configure an allowed list on one switch in the fabric, we recommend you configure the same list in all other switches in the fabric to ensure consistency.

An allowed domain ID list must satisfy the following conditions:

- If this switch is a principal switch, all the currently assigned domain IDs must be in the allowed list.
- If this switch is a subordinate switch, the local runtime domain ID must be in the allowed list.
- The locally configured domain ID of the switch must be in the allowed list.
- The intersection of the assigned domain IDs with other already configured domain ID lists must not be empty.

To configure the allowed domain ID list, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fcdomain allowed 50-110 vsan 4 | Configures the list to allow switches with the domain ID 50 through 110 in VSAN 4. |
| | switch(config)# no fcdomain allowed 50-110 vsan 5 | Reverts to the factory default of allowing domain IDs from 1 through 239 in VSAN 5. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Merged Stable Fabrics

By default, the **auto-reconfigure** option is disabled. When you join two switches belonging to two different stable fabrics that have overlapping domains, the following cases apply:

- If the **auto-reconfigure** option is enabled on both switches, a disruptive reconfiguration phase is started.
- If the **auto-reconfigure** option is disabled on either or both switches, the links between the two switches become isolated.

The **auto-reconfigure** option takes immediate effect at runtime. You do not need to restart the fcdomain. If a domain is currently isolated due to domain overlap, and you later enable the **auto-reconfigure** option on both switches, the fabric continues to be isolated—if you enabled the **auto-reconfigure** option on both switches before connecting the fabric, a disruptive reconfiguration (RCF) will occur. A disruptive reconfiguration may affect data traffic. You can nondisruptively reconfigure the fcdomain by changing the configured domains on the overlapping links and getting rid of the domain overlap.

To enable automatic reconfiguration in a specific VSAN (or range of VSANs), follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fcdomain auto-reconfigure vsan 10 | Enables the automatic reconfiguration option in VSAN 10. |
| | switch(config)# no fcdomain auto-reconfigure 69 | Disables the automatic reconfiguration option and reverts it to the factory default in VSAN 69. |

Contiguous Domain Assignments

By default, the **contiguous-allocation** option is disabled. When a subordinate switch requests the principal switch for two or more domains and the domains are not contiguous, the following cases apply:

- If the **contiguous-allocation** option is enabled in the principal switch, the principal switch locates contiguous domains and assigns them to the subordinate switches. If contiguous domains are not available, the SAN-OS software rejects this request.
- If the **contiguous-allocation** option is disabled in the principal switch, the principal switch assigns the available domains to the subordinate switch.

The **contiguous-allocation** option takes immediate effect at runtime—you do not need to restart the fcdomain.

To enable contiguous domains in a specific VSAN (or a range of VSANs), follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fcdomain contiguous-allocation vsan 81-83 | Enables the contiguous allocation option in VSAN 81 through 83. |
| | switch(config)# no fcdomain contiguous-allocation vsan 1030 | Disables the contiguous allocation option and reverts it to the factory default in VSAN 1030. |

Send documentation comments to mdsfeedback-doc@cisco.com.

fcdomain Initiation

By default, the fcdomain feature is enabled on each switch. If you disable the fcdomain feature in a switch, that switch can no longer participate with other switches in the fabric. The fcdomain configuration is applied to runtime through a disruptive restart.

Use the **no fcdomain** command to disable the fcdomain feature.

To disable fcdomains in a single VSAN or a range of VSANs, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# no fcdomain vsan 7-200 | Disables the fcdomain configuration in VSAN 7 through 200. |
| | switch(config)# fcdomain vsan 2008 | Enables the fcdomain configuration in VSAN 2008. |

Fabric Name

By default the configured fabric name is 20:01:00:05:30:00:28:df.

- When the fcdomain feature is disabled, the runtime fabric name is the same as the configured fabric name.
- When the fcdomain feature is enabled, the runtime fabric name is the same as the principal switch's WWN.

The fabric name is applied to runtime through a disruptive restart when the fcdomain is configured as disabled (see the [“Domain Restart” section on page 31-3](#)).

To set the fabric name value for a disabled fcdomain, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fcdomain fabric-name 20:1:ac:16:5e:0:21:01 vsan 3 | Assigns the configured fabric name value in VSAN 3. |
| | switch(config)# no fcdomain fabric-name 20:1:ac:16:5e:0:21:01 vsan 3010 | Changes the fabric name value to the factory default (20:01:00:05:30:00:28:df) in VSAN 3010. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Incoming RCFs

The **rcf-reject** option is configured on a per-interface, per-VSAN basis. By default, the **rcf-reject** option is disabled (that is, RCF request frames are not automatically rejected).

The **rcf-reject** option takes immediate effect to runtime through a disruptive restart (see the “[Domain Restart](#)” section on page 31-3).

To stop incoming RCF request frames, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/1 switch(config-if)# | Configures the specified interface. |
| Step 3 | switch(config-if)# fcdomain rcf-reject vsan 1 | Enables the RCF filter on the specified interface in VSAN 1. |
| | switch(config-if)# no fcdomain rcf-reject vsan 1 | Disables the RCF filter on the specified interface in VSAN 1. |

Persistent FC IDs

As of Cisco SAN-OS Release 2.0(1b), persistent FC IDs are enabled by default. This change prevents FCIDs from being changed after a reboot. You can disable this option for each VSAN.

When an N or NL port logs into a Cisco MDS 9000 Family switch, it is assigned a FC ID. By default, the persistent FC ID feature is enabled. If this feature is disabled, the following consequences apply:

- An N or NL port logs into a Cisco MDS 9000 Family switch. The WWN of the requesting N or NL port and the assigned FC ID, are retained and stored in a volatile cache. The contents of this volatile cache are not saved across reboots.
- The switch is designed to preserve the binding FC ID to the WWN on a best-effort basis. For example, if one N port disconnects from the switch and its FC ID is requested by another device, this request is granted, and the WWN with the initial FC ID association is released.
- The volatile cache stores up to 4000 entries of WWN to FC ID binding. If this cache is full, a new (more recent) entry overwrites the oldest entry in the cache. In this case, the corresponding WWN to FC ID association for the oldest entry is lost.
- The switch connection behavior differs between N ports and NL ports:
 - N ports receive the same FC IDs if disconnected and reconnected to any port within the same switch (as long as it belongs to the same VSAN).
 - NL ports receive the same FC IDs only if connected back to the same port on the switch to which they were originally connected.

Send documentation comments to mdsfeedback-doc@cisco.com.

If this feature remains enabled, the following consequences apply:

- The currently *in use* FC IDs in the fcdomain are saved across reboots.
- The fcdomain automatically populates the database with dynamic entries that the switch has learned about after a device (host or disk) is plugged into a port interface.



Note If you connect to the switch from an AIX or HP-UX host, be sure to enable the persistent FC ID feature in the VSAN that connects these hosts.

A persistent FC ID assigned to an F port can be moved across interfaces and can continue to maintain the same persistent FC ID.



Note Persistent FC IDs with loop-attached devices (FL ports) need to remain connected to the same port in which they were configured.

To enable the persistent FC ID feature, follow these steps:

| | Command | Purpose |
|---------------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# fcdomain fcid persistent vsan 1000 FCID(s) persistent feature is enabled. | Activates (default) persistency of FC IDs in VSAN 1000. |
| | switch(config)# no fcdomain fcid persistent vsan 20 | Disables the FC ID persistency feature in VSAN 20. |

Persistent FC IDs Manual Configuration

When the persistent FC ID feature is enabled, you can enter the persistent FC ID submode and add static or dynamic entries in the FC ID database. By default, all added entries are static. Persistent FC IDs are configured on a per-VSAN basis. Follow these requirements to manually configure a persistent FC ID:

- Ensure that the persistent FC ID feature is enabled in the required VSAN.
- Ensure that the required VSAN is an active VSAN—persistent FC IDs can only be configured on active VSANs.
- Verify that the domain part of the FC ID is the same as the runtime domain ID in the required VSAN. If the software detects a domain mismatch, the command is rejected.
- Verify that the port field of the FC ID is 0 (zero) when configuring an area.



Note FICON uses a different scheme for allocating FC IDs based in the front panel port number. This scheme takes precedence over FC ID persistence in FICON VSANs.

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure persistent FC IDs, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fcdomain fcid database switch(config-fcid-db)# | Enters FC ID database configuration submode. |
| Step 3 | switch(config-fcid-db)# vsan 1000 wwn 33:e8:00:05:30:00:16:df fcid 0x070128 | Configures a device WWN (33:e8:00:05:30:00:16:df) with the FC ID 0x070128 in VSAN 1000. |
| | switch(config-fcid-db)# vsan 1000 wwn 11:22:11:22:33:44:33:44 fcid 0x070123 dynamic | Configures a device WWN (11:22:11:22:33:44:33:44) with the FC ID 0x070123 in VSAN 1000 in dynamic mode. |
| | switch(config-fcid-db)# vsan 1000 wwn 11:22:11:22:33:44:33:44 fcid 0x070100 area | Configures a device WWN (11:22:11:22:33:44:33:44) with the FC IDs 0x070100 through 0x701FF in VSAN 1000. |
| | | Note To secure the entire area for this fcdomain, assign 00 as the last two characters of the FC ID. |

Unique Area FC IDs for Some HBAs



Note

Only read this section if the HBA port and the storage port are connected to the same switch.

Some HBA ports require a different area ID than storage ports when they are both connected to the same switch. For example, if the storage port FC ID is 0x6f7704, the area for this port is 77. In this case, the HBA port's area can be anything other than 77. The HBA port's FC ID must be manually configured to be different from the storage port's FC ID.

Switches in the Cisco MDS 9000 Family facilitate this requirement with the FC ID persistence feature. You can use this feature to preassign an FC ID with a different area to either the storage port or the HBA port. The procedure in this example uses a switch domain of 111(6f hex). The HBA port connects to interface fc1/9 and the storage port connects to interface fc 1/10 in the same switch.

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure a different area ID for the HBA port, follow these steps:

- Step 1** Obtain the Port WWN (Port Name field) ID of the HBA using the **show flogi database** command).

```
switch# show flogi database
```

| INTERFACE | VSAN | FCID | PORT NAME | NODE NAME |
|-----------|------|----------|-------------------------|-------------------------|
| fc1/9 | 3 | 0x6f7703 | 50:05:08:b2:00:71:c8:c2 | 50:05:08:b2:00:71:c8:c0 |
| fc1/10 | 3 | 0x6f7704 | 50:06:0e:80:03:29:61:0f | 50:06:0e:80:03:29:61:0f |



Note Both FC IDs in this setup have the same area 77 assignment.

- Step 2** Shut down the HBA interface in the MDS switch.

```
switch# conf t
switch(config)# interface fc1/9
switch(config-if)# shutdown
switch(config-if)# end
switch#
```

- Step 3** Verify that the FC ID feature is enabled using the **show fcdomain vsan** command.

```
switch# show fcdomain vsan 1
...
Local switch configuration information:
    State: Enabled
    FCID persistence: Disabled
```

If this feature is disabled, continue with this procedure to enable the FC ID persistence.

If this feature is already enabled, skip to [Step 5](#).

- Step 4** Enable the FC ID persistence feature in the Cisco MDS switch.

```
switch# conf t
switch(config)# fcdomain fcid persistent vsan 1
switch(config)# end
switch#
```

- Step 5** Assign a new FC ID with a different area allocation. In this example, we replace 77 with *ee*.

```
switch# conf t
switch(config)# fcdomain fcid database
switch(config-fcid-db)# vsan 3 wwn 50:05:08:b2:00:71:c8:c2 fcid 0x6fee00 area
```

- Step 6** Enable the HBA interface in the Cisco MDS switch.

```
switch# conf t
switch(config)# interface fc1/9
switch(config-if)# no shutdown
switch(config-if)# end
switch#
```


Send documentation comments to mdsfeedback-doc@cisco.com.

Step 7 Verify the pWWN ID of the HBA using the **show flogi database** command.

```
switch# show flogi database
```

```
-----
INTERFACE   VSAN      FCID      PORT NAME      NODE NAME
-----
fc1/9       3         0x6fee00   50:05:08:b2:00:71:c8:c2   50:05:08:b2:00:71:c8:c0
fc1/10      3         0x6f7704   50:06:0e:80:03:29:61:0f   50:06:0e:80:03:29:61:0f
```



Note Both FC IDs now have different area assignments.

Persistent FC ID Selective Purging

Persistent FC IDs can be purged selectively. Static entries and FC IDs currently in use cannot be deleted. [Table 31-1](#) identifies the FC ID entries that are deleted or retained when persistent FC IDs are purged.

Table 31-1 *Purged FC IDs*

| Persistent FC ID state | Persistent Usage State | Action |
|------------------------|------------------------|-------------|
| Static | In use | Not deleted |
| Static | Not in use | Not deleted |
| Dynamic | In use | Not deleted |
| Dynamic | Not in use | Deleted |

Use the **purge fcdomain** command to remove entries which dynamic and unused (see [Table 31-1](#)).

To purge persistent FC IDs, follow this step:

| Command | Purpose |
|---|---|
| Step 1 switch# purge fcdomain fcid vsan 4 | Purges all dynamic and unused FC IDs in VSAN 4. |
| switch# purge fcdomain fcid vsan 3-5 | Purges dynamic and unused FC IDs in VSAN 3, 4, and 5. |

Displaying fcdomain Information

Use the **show fcdomain** command to display global information about fcdomain configurations. See [Example 31-1](#).



Note In [Example 31-1](#), the fcdomain feature is disabled. Consequently, the runtime fabric name is the same as the configured fabric name.

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 31-1 *Displays the Global fcdomain Information*

```
switch# show fcdomain vsan 2
The local switch is the Principal Switch.

Local switch run time information:
  State: Stable
  Local switch WWN:      20:01:00:0b:46:79:ef:41
  Running fabric name: 20:01:00:0b:46:79:ef:41
  Running priority: 128
  Current domain ID: 0xed(237)

Local switch configuration information:
  State: Enabled
  FCID persistence: Disabled
  Auto-reconfiguration: Disabled
  Contiguous-allocation: Disabled
  Configured fabric name: 20:01:00:05:30:00:28:df
  Configured priority: 128
  Configured domain ID: 0x00(0) (preferred)

Principal switch run time information:
  Running priority: 128

No interfaces available.
```

Use the **show fcdomain domain-list** command to display the list of domain IDs of all switches belonging to a specified VSAN. This list provides the WWN of the switches owning each domain ID.

[Example 31-2](#) shows the following:

- A switch with WWN of 20:01:00:05:30:00:47:df is the principal switch and has domain 200.
- A switch with WWN of 20:01:00:0d:ec:08:60:c1 is the local switch (the one where you typed the CLI command to show the domain-list) and has domain 99.
- The IVR manager obtained virtual domain 97 using 20:01:00:05:30:00:47:df as the WWN for a virtual switch.

Example 31-2 *Displays the fcdomain Lists*

```
switch# show fcdomain domain-list vsan 76

Number of domains: 3
Domain ID          WWN
-----
0xc8(200)          20:01:00:05:30:00:47:df [Principal]
0x63(99)            20:01:00:0d:ec:08:60:c1 [Local]
0x61(97)            50:00:53:0f:ff:f0:10:06 [Virtual (IVR)]
```

Use the **show fcdomain allowed vsan** command to display the list of allowed domain IDs configured on this switch. See [Example 31-3](#).

Example 31-3 *Displays the Allowed Domain ID Lists*

```
switch# show fcdomain allowed vsan 1
Assigned or unallowed domain IDs: 1-96,100,111-239.
[Interoperability Mode 1] allowed domain IDs: 97-127.
[User] configured allowed domain IDs: 50-110.
```

Send documentation comments to mdsfeedback-doc@cisco.com.



Tip

Ensure that the requested domain ID passes the Cisco SAN-OS software checks, if **interop 1** mode is required in this switch.

Use the **show fcdomain fcid persistent** command to display all existing, persistent FC IDs for a specified VSAN. You can also specify the **unused** option to view only persistent FC IDs that are still not in use. See Examples 31-4 and 31-5.

Example 31-4 Displays Persistent FC IDs in a Specified VSAN

```
switch# show fcdomain fcid persistent vsan 1000
Total entries 2.
```

Persistent FCIDs table contents:

| VSAN | WWN | FCID | Mask | Used | Assignment |
|------|-------------------------|----------|-------------|------|------------|
| 1000 | 11:11:22:22:11:11:12:23 | 0x700101 | SINGLE FCID | NO | STATIC |
| 1000 | 44:44:33:33:22:22:11:11 | 0x701000 | ENTIRE AREA | NO | DYNAMIC |

Example 31-5 Displays All Persistent FC IDs in the fcdomain

```
switch# show fcdomain fcid persistent
Total entries 2.
```

Persistent FCIDs table contents:

| VSAN | WWN | FCID | Mask | Used | Assignment |
|------|-------------------------|----------|-------------|------|------------|
| 1000 | 11:11:22:22:11:11:22:22 | 0x700501 | SINGLE FCID | NO | STATIC |
| 1003 | 44:44:33:33:22:22:11:11 | 0x781000 | ENTIRE AREA | YES | DYNAMIC |

Use the **show fcdomain statistics** command to display frame and other fcdomain statistics for a specified VSAN or PortChannel. See Example 31-6 and Example 31-7.

Example 31-6 Displays fcdomain Statistics for a Specified VSAN

```
switch# show fcdomain statistics vsan 1
VSAN Statistics
  Number of Principal Switch Selections: 5
  Number of times Local Switch was Principal: 0
  Number of 'Build Fabric's: 3
  Number of 'Fabric Reconfigurations': 0
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 31-7 Displays fcdomain Statistics for a Specified PortChannel

```
switch# show fcdomain statistics interface port-channel 10 vsan 1
Interface Statistics:
      Transmitted      Received
      -----
      EFPs           13           9
      DIAs            7           7
      RDIs            0           0
      ACCs           21          25
      RJTs            1           1
      BFs             2           2
      RCFs            4           4
      Error            0           0
      Total           48          48
Total Retries: 0
Total Frames: 96
      -----
```

Use the **show fcdomain address-allocation** command to display FC ID allocation statistics including a list of assigned and free FC IDs. See [Example 31-8](#).

Example 31-8 Displays FC ID Information

```
switch# show fcdomain address-allocation vsan 1
Free FCIDs: 0x020000 to 0x02fdff
            0x02ff00 to 0x02fffe

Assigned FCIDs: 0x02fe00 to 0x02feff
                0x02ffff

Reserved FCIDs: 0x020100 to 0x02f0ff
                0x02fe00 to 0x02feff
                0x02ffff

Number free FCIDs: 65279
Number assigned FCIDs: 257
Number reserved FCIDs: 61697
```

Use the **show fcdomain address-allocation cache** command to display the valid address allocation cache. The cache is used by the principal switch to reassign the FC IDs for a device (disk or host) that exited and reentered the fabric. In the cache content, VSAN refers to the VSAN that contains the device, WWN refers to the device that owned the FC IDs, and mask refers to a single or entire area of FC IDs. See [Example 31-9](#).

Example 31-9 Displays Address Allocation Information

```
switch# show fcdomain address-allocation cache
Cache content:
line#   VSAN      WWN                      FCID      mask
-----
  1.    12      21:00:00:e0:8b:08:a2:21  0xef0400  ENTIRE AREA
  2.     6      50:06:04:82:c3:a1:2f:5c  0xef0002  SINGLE FCID
  3.     8      20:4e:00:05:30:00:24:5e  0xef0300  ENTIRE AREA
  4.     8      50:06:04:82:c3:a1:2f:52  0xef0001  SINGLE FCID
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Default Settings

Table 31-2 lists the default settings for all fcdomain parameters.

Table 31-2 **Default fcdomain Parameters**

| Parameters | Default |
|-------------------------------------|--|
| fcdomain feature | Enabled. |
| Configured domain ID | 0 (zero). |
| Configured domain | Preferred. |
| auto-reconfigure option | Disabled. |
| contiguous-allocation option | Disabled. |
| Priority | 128. |
| Allowed list | 1 to 239. |
| Fabric name | 20:01:00:05:30:00:28:df. |
| rcf-reject | Disabled. |
| Persistent FC ID | Enabled (as of Release 2.0(1b) this is only configurable on a per-VSAN basis). |

Send documentation comments to mdsfeedback-doc@cisco.com.



Configuring Traffic Management

Fibre Channel Congestion Control (FCC) is a Cisco proprietary flow control mechanism that alleviates congestion on Fibre Channel networks.

Quality of service (QoS) offers the following advantages:

- Provides relative bandwidth guarantee to application traffic.
- Controls latency experienced by application traffic.
- Prioritizes one application over another (for example, prioritizing transactional traffic over bulk traffic) through bandwidth and latency differentiation.

This chapter provides details on the QoS and FCC features provided in all switches. It includes the following sections:


- [FCC, page 32-2](#)
- [QoS, page 32-3](#)
- [Control Traffic, page 32-4](#)
- [Data Traffic, page 32-4](#)
- [Ingress Port Rate Limiting, page 32-13](#)
- [Default Settings, page 32-14](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

FCC

FCC reduces the congestion in the fabric without interfering with the standard Fibre Channel protocols. The FCC protocol increases the granularity and the scale of congestion control applied to any class of traffic (see [Figure 32-1](#)).

Figure 32-1 FCC Mechanisms



Edge quench congestion control provides feedback to the source about the rate at which frames should be injected into the network (frame intervals).

FCC Process

When a node in the network detects congestion for an output port, it generates an edge quench message. These frames are identified by the Fibre Channel destination ID (DID) and the source ID. A switch from other vendors simply forwards these frames.

Any receiving switch in the Cisco MDS 9000 Family handles frames in one of these ways:

- It forwards the frame.
- It limits the rate of the frame flow in the congested port.

The behavior of the flow control mechanism differs based on the Fibre Channel DID:

- If the Fibre Channel DID is directly connected to one of the switch ports, the input rate limit is applied to that port.
- If the destination of the edge quench frame is a Cisco domain or the next hop is a Cisco MDS 9000 Family switch, the frame is forwarded.
- If neither of these mechanisms is true, then the frame is processed in the port going towards the FC DID.

All switches (including the edge switch) along the congested path process path quench frames. However, only the edge switch processes edge quench frames.

Send documentation comments to mdsfeedback-doc@cisco.com.

Enabling FCC

By default, the FCC protocol is disabled. FCC can only be enabled for the entire switch.



Tip

If you enable FCC, be sure to enable it in all switches in the fabric.

To enable or disable the FCC feature, follow these steps:

| | Command | Purpose |
|--------|-------------------------------|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# fcc | Enables FCC in this switch. |
| | switch(config)# no fcc | Disables FCC in this switch (default). |

Assigning FCC Priority

To assign FCC priority, follow these steps:

| | Command | Purpose |
|--------|---------------------------------------|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# fcc priority 2 | Defines the FCC priority threshold to have a priority of 2—0 is the lowest priority and 7 is the highest priority. |

Displaying FCC

Use the **show fcc** command to view FCC settings (see [Example 32-1](#)).

Example 32-1 Displays Configured FCC Information

```
switch# show fcc
fcc is disabled
fcc is applied to frames with priority up to 4
```

QoS

QoS implementation in the Cisco MDS 9000 Family follows the differentiated services (DiffServ) model. The DiffServ standard is defined in RFCs 2474 and 2475.

All switches support the following types of traffic:

- [Control Traffic](#), page 32-4
- [QoS](#), page 32-3

Send documentation comments to mdsfeedback-doc@cisco.com.

Control Traffic

- The Cisco MDS 9000 Family supports QoS for internally and externally generated control traffic. Within a switch, control traffic is sourced to the supervisor module and is treated as a high priority frame. A high priority status provides absolute priority over all other traffic and is assigned in the following cases:
- Internally generated time-critical control traffic (mostly Class F frames).
 - Externally generated time-critical control traffic entering a switch in the Cisco MDS 9000 Family from a another vendor’s switch. High priority frames originating from other vendor switches are marked as high priority as they enter a switch in the Cisco MDS 9000 Family.

Disabling Control Traffic

By default, the QoS feature for certain critical control traffic is enabled. These critical control frames are assigned the highest (absolute) priority.



Tip

We do not recommend disabling this feature as all critical control traffic is automatically assigned the lowest priority once you issue this command.

To disable the high priority assignment for control traffic, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# no qos control priority 0 | Enables the control traffic QoS feature. |
| | switch(config)# qos control priority 0 | Disables the control traffic QoS feature. |

Displaying Control Traffic Information

Use the **show qos statistics** command to view the current state of the QoS configuration for critical control traffic. This command displays the current QoS settings along with the number of frames marked high priority. The count is only for debugging purposes and cannot be configured (see [Example 32-2](#)).

Example 32-2 *Displays Current QoS Settings*

```
switch# show qos statistics
Total number of FC frames transmitted from the Supervisor= 15767
Number of highest-priority FC frames transmitted           = 8224
Current priority of FC control frames = 0      (0 = lowest; 7 = highest)
```

Data Traffic

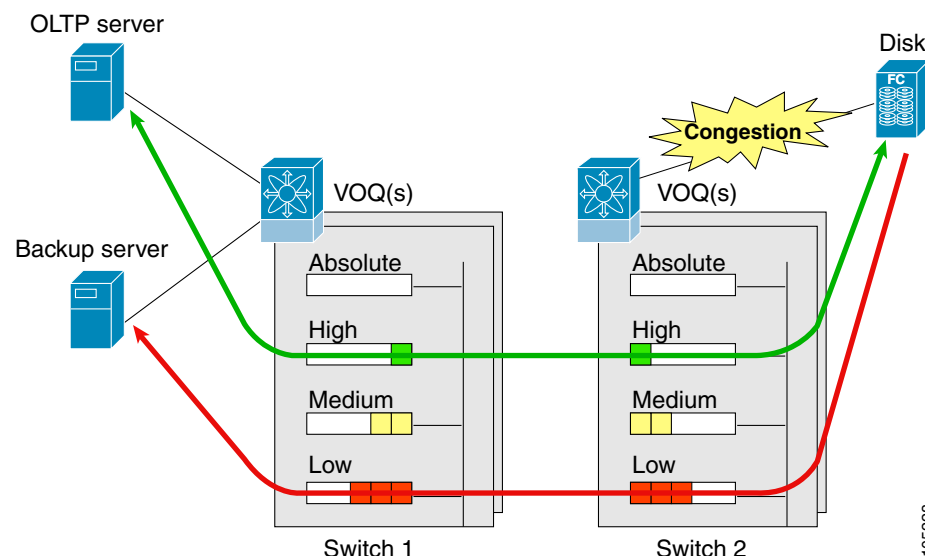
Online transaction processing, which is a low volume, latency sensitive application, requires quick access to requested information. Backup processing application require high bandwidth but are not sensitive to latency. In a network that does not support service differentiation, all traffic is treated

Send documentation comments to mdsfeedback-doc@cisco.com.

identically—they experience similar latency and are allocated similar bandwidths. The QoS feature in the Cisco MDS 9000 Family switches provides these guarantees as of Cisco MDS SAN-OS Release 1.3(1).

Earlier versions of the Cisco MDS SAN-OS software only differentiated traffic priority based on control traffic. Cisco MDS SAN-OS Release 1.3(1) enables you to take full advantage of the QoS capabilities. Data traffic can be prioritized in distinct levels of service differentiation: low, medium, or high priority. You can apply QoS to ensure that Fibre Channel data traffic for your latency-sensitive applications receive higher priority over throughput-intensive applications such as data warehousing (see Figure 32-2).

Figure 32-2 Prioritizing Data Traffic



In Figure 32-2, the OLTP traffic arriving at Switch 1 is marked with a high priority level of through classification (class map) and marking (policy map). Similarly, the backup traffic is marked with a low priority level. The traffic is sent to the corresponding priority queue within a virtual output queue (VOQ).

A deficit weighted round robin (DWRR) scheduler configured in the first switch ensures that high priority traffic is treated better than low priority traffic. For example, DWRR weights of 70:20:10 implies that the high priority queue is serviced at 7 times the rate of the low priority queue. This guarantees lower delays and higher bandwidths to high priority traffic if congestion sets in. A similar configuration in the second switch ensures the same traffic treatment in the other direction.

If the ISL is congested when the OLTP server sends a request, the request is queued in the high priority queue and is serviced almost immediately as the high priority queue is not congested. The scheduler assigns it priority over the backup traffic in the low priority queue.



Note

When the high priority queue does not have traffic flowing through, the low priority queue uses all the bandwidth and is not restricted to the configured value.

A similar occurrence in Switch 2 sends a response to the transaction request. The round trip delay experienced by the OLTP server is independent of the volume of low priority traffic or the ISL congestion. The backup traffic uses the available ISL bandwidth when it is not used by the OLTP traffic.

Send documentation comments to mdsfeedback-doc@cisco.com.



Tip

To achieve this traffic differentiation, be sure to enable FCC (see the “[Enabling FCC](#)” section on [page 32-3](#)).

VSAN Versus Zone-Based QoS

While you can configure both zone-based QoS and VSAN-based QoS configurations in the same switch, both configurations have significant differences. [Table 32-1](#) highlights the differences between configuring QoS priorities based on VSANs versus zones.

Table 32-1 QoS Configuration Differences

| VSAN-Based QoS | Zone-Based QoS |
|--|--|
| If you configure the active zone set on a given VSAN and also configure QoS parameters in any of the member zones, you will not be able to associate the policy map with the VSAN. | You cannot activate a zone set on a VSAN which already has a policy map associated. |
| If the same flow is present in two class maps associated to a policy map, the QoS value of the class map attached first takes effect. | If the same flow is present in two zones in a given zoneset with different QoS values, the higher QoS value is considered. |
| — | During a zone merge, if the Cisco SAN-OS software detects a mismatch for the QoS parameter, the link is isolated. |
| Takes effect even if QoS is disabled. | Takes effect only when QoS is enabled. |

See the “[Zone-Based Traffic Priority](#)” section on [page 15-15](#) for details on configuring a zone-based QoS policy.

Configuring Data Traffic

To configure QoS, follow these steps.

- Step 1** Enable the QoS feature.
- Step 2** Create and define class maps.
- Step 3** Define service policies.
- Step 4** Apply the configuration.

QoS Initiation for Data Traffic

By default, the QoS data traffic feature is disabled for data traffic. To configure QoS for data traffic, you must first enable the data traffic feature in the switch.

Send documentation comments to mdsfeedback-doc@cisco.com.



Tip

QoS is supported in interoperability mode—its effectiveness depends on the location of Cisco MDS 9000 Family switches in the fabric relative to the location of the source or destination of the prioritized devices.

To enable the QoS data traffic feature, follow these steps:

| | Command | Purpose |
|---------------|--------------------------------------|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# qos enable | Enables QoS. You can now configure data traffic parameters. |
| | switch(config)# no qos enable | Removes the currently applied QoS configuration and disables QoS. You can no longer configure data traffic parameters. |

Class Map Creation

Use the class map feature to create and define a traffic class with match criteria to identify traffic belonging to that class. The class map name is restricted to 63 alphanumeric characters and defaults to the **match-all** option. Flow-based traffic uses one of the following values:

- **WWN**—The source WWN or the destination WWN.
- **Fibre Channel ID (FC ID)** —The source ID (SID) or the destination ID (DID). The possible values for mask are FFFFFFFF (the entire FC ID is used—this is the default), FFFF00 (only domain and area FC ID is used), or FF0000 (only domain FC ID is used).



Note

A SID or DID of 0x000000 is not allowed.

- **Source interface**—The ingress interface.



Tip

The order of entries to be matched within a class map is not significant.

Use the **class-map** command to create and define a traffic class with match criteria to identify traffic belonging to that class. Define each match criterion with one match statement from the class map configuration (`switch(config-cmap)`) mode.

- Use the **source-wwn** option to specify the source WWN or the **destination-wwn** option to specify the destination WWN.
- Use the **source-address** option to specify the source ID (SID) or the **destination-address** option to specify the destination ID (DID).
- Use the **input-interface** option to specify the ingress interface.
- Use the **destination-device-alias** option to specify the distributed device alias.

Send documentation comments to mdsfeedback-doc@cisco.com.

To create a class map, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | <code>switch(config)# qos class-map MyClass</code> <code>switch(config-cmap)#</code> | Creates a class map called MyClass and places you in the class-map submode to match all criteria specified for this class. |
| | <code>switch(config)# qos class-map MyClass</code> <code>match-all</code> <code>switch(config-cmap)#</code> | Specifies a logical AND operator for all matching statements in this class. If a frame matches all (default) configured criteria, it qualifies for this class. This is the default. |
| | <code>switch(config)# qos class-map MyClass</code> <code>match-any</code> <code>switch(config-cmap)#</code> | Specifies a logical OR operator for all matching statements in this class. If a frame matches any one configured criteria, it qualifies for this class. |
| Step 2 | <code>switch(config-cmap)# match</code> <code>destination-address 0x12ee00</code> | Specifies a destination address match for frames with the specified destination FC ID. |
| | <code>switch(config-cmap)# match source-address</code> <code>0x6d1090 mask 0xFFFFF</code> | Specifies a source address and mask match for frames with the specified source FC ID. |
| Step 3 | <code>switch(config-cmap)# match destination-wwn</code> <code>20:01:00:05:30:00:28:df</code> | Specifies a destination WWN to match frames. |
| | <code>switch(config-cmap)# match source-wwn</code> <code>23:15:00:05:30:00:2a:1f</code> | Specifies a source WWN to match frames. |
| Step 4 | <code>switch(config-cmap)# match</code> <code>destination-device-alias DocDeviceAlias</code> | Specifies a destination device alias to match frames. |
| Step 5 | <code>switch(config-cmap)# match input-interface fc</code> <code>2/1</code> | Specifies a source interface to match frames. |
| Step 6 | <code>switch(config-cmap)# no match input-interface</code> <code>fc 3/5</code> | Removes a match based on the specified source interface. |

Service Policy Definition

Service policies are specified using policy maps. Policy maps provide an ordered mapping of class maps to service levels. You can specify multiple class maps within a policy map, and map a class map to a high, medium, or low service level. The default priority is low. The policy map name is restricted to 63 alphanumeric characters.

As an alternative, you can map a class map to a differentiated services code point (DSCP). The DSCP is an indicator of the service level for a specified frame. The DSCP value ranges from 0 to 63, and the default is 0. A DSCP value of 46 is disallowed.

The order of the class maps within a policy map is important to determine the order in which the frame is compared to class maps. The first matching class map has the corresponding priority marked in the frame.



Note

Refer to <http://www.cisco.com/warp/public/105/dscpvalues.html#dscpandassuredforwardingclasses> for further information on implementing QoS DSCP values.



Note

Class maps are processed in the order in which they are configured in each policy map.

Send documentation comments to mdsfeedback-doc@cisco.com.

Use the **policy-map** option to specify the class of service.

To specify a service policy, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | <code>switch(config)# qos policy-map MyPolicy</code> <code>switch(config-pmap)#</code> | Creates a policy map called MyPolicy and places you in the policy-map submode. |
| | <code>switch(config)# no qos policy-map OldPolicy</code> <code>switch(config)#</code> | Deletes the policy map called OldPolicy and places you in the policy-map submode. |
| Step 2 | <code>switch(config-pmap)# class MyClass</code> <code>switch(config-pmap-c)#</code> | Specifies the name of a predefined class and places you at the policy-map submode for that class. |
| | <code>switch(config-pmap)# no class OldClass</code> | Removes the class map called OldClass from the policy map. |
| Step 3 | <code>switch(config-pmap-c)# priority high</code> | Specifies the priority to be given for each frame matching this class. |
| | <code>switch(config-pmap-c)# no priority high</code> | Deletes a previously assigned priority and reverts to the default value of low. |
| Step 4 | <code>switch(config-pmap-c)# dscp 2</code> | Specifies the DSCP value to mark each frame matching this class. |
| | <code>switch(config-pmap-c)# no dscp 60</code> | Deletes a previously assigned DSCP value and reverts to the factory default of 0. |

Service Policy Enforcement

When you have configured a QoS data traffic policy, you must enforce the data traffic configuration by applying that policy to the required VSAN(s). If you do not apply the policy to a VSAN, the data traffic configuration is not enforced. You can only apply one policy map to a VSAN.



Note

You can apply the same policy to a range of VSANs.

To apply a service policy, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | <code>switch(config)# qos service policy MyPolicy</code> <code>vsan 3</code> | Applies a configured policy to VSAN 3. |
| | <code>switch(config)# no qos service policy OldPolicy</code> <code>vsan 7</code> | Deletes a configured policy that was applied to VSAN 7. |

DWRR Traffic Scheduler

The Cisco SAN-OS software supports four scheduling queues:

- Strict priority queues are queues that are serviced in preference to other queues—it is always serviced if there is a frame queued in it regardless of the state of the other queues.
- QoS assigns all other traffic to the DWRR scheduling high, medium, and low priority traffic queues.

Send documentation comments to mdsfeedback-doc@cisco.com.

The DWRR scheduler services the queues in the ratio of the configured weights. Higher weights translate to proportionally higher bandwidth and lower latency. The default weights are 50 for the high queue, 30 for the medium queue, and 20 for the low queue. Decreasing order of queue weights is mandated to ensure the higher priority queues have a higher service level, though the ratio of the configured weights can vary (for example, one can configure 70:30:5 or 60:50:10 but not 50:70:10).

Use the **qos dwrr-q** command to associate a weight with a DWRR queue. Use the **dwrr-q high** option to schedule high priority traffic, the **dwrr-q medium** option to schedule medium priority traffic, and the **dwrr-q low** option to schedule low priority traffic.

To associate a weight with a DWRR queue, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | <code>switch(config)# qos dwrr-q high weight 10</code> | Associates a relative weight (10) to a specified queue (default queue). |
| | <code>switch(config)# no qos dwrr-q low weight 51</code> | Restores the default weight of 20. |

Displaying Data Traffic Information

The **show qos** commands display the current QoS settings for data traffic (see Examples 32-3 to 32-11).

Example 32-3 Displays the Contents of all Class Maps

```
switch# show qos class-map
qos class-map MyClass match-any
  match destination-wwn 20:01:00:05:30:00:28:df
  match source-wwn 23:15:00:05:30:00:2a:1f
  match input-interface fc2/1
qos class-map Class2 match-all
  match input-interface fc2/14
qos class-map Class3 match-all
  match source-wwn 20:01:00:05:30:00:2a:1f
```

Example 32-4 Displays the Contents of a Specified Class Map

```
switch# show qos class-map name MyClass
qos class-map MyClass match-any
  match destination-wwn 20:01:00:05:30:00:28:df
  match source-wwn 23:15:00:05:30:00:2a:1f
  match input-interface fc2/1
```

Example 32-5 Displays All Configured Policy Maps

```
switch# show qos policy-map
qos policy-map MyPolicy
  class MyClass
    priority medium
qos policy-map Policy1
  class Class2
    priority low
```


Send documentation comments to mdsfeedback-doc@cisco.com.

Example 32-6 Displays a Specified Policy Map

```
switch# show qos policy-map name MyPolicy
qos policy-map MyPolicy
  class MyClass
    priority medium
```

Example 32-7 Displays Scheduled DWRR Configurations

```
switch# show qos dwrr
qos dwrr-q high weight 50
qos dwrr-q medium weight 30
qos dwrr-q low weight 20
```

Example 32-8 Displays All Applied Policy Maps

```
switch# show qos service policy
qos service policy MyPolicy vsan 1
qos service policy Policy1 vsan 4
```

Example 32-9 Displays the Policy Map Associated with a Specified VSAN

```
switch# show qos service policy vsan 1
qos policy-map pmap1
  class cmap1
    priority medium
  class cmap2
    priority high
```

Example 32-10 Displays the Class Map Associated with a Specified Interface

```
switch# show qos service policy interface fc3/10
qos policy-map pmap1
  class cmap3
    priority high
  class cmap4
    priority low
```

Example 32-11 Displays QoS Statistics

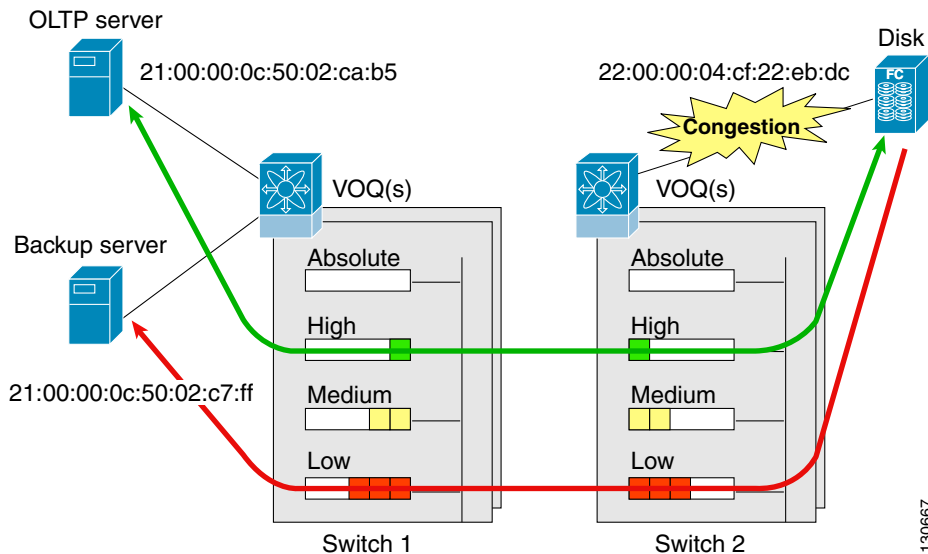
```
switch# show qos statistics
Total number of FC frames transmitted from the Supervisor= 301431
Number of highest-priority FC frames transmitted           = 137679
Current priority of FC control frames = 7      (0 = lowest; 7 = highest)
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example Configuration

This section describes a configuration example for the application illustrated in Figure 32-3.

Figure 32-3 Example Application for Traffic Prioritization



Both the OLTP server and the backup server are accessing the disk. The backup server is writing large amounts of data to the disk. This data does not require specific service guarantees. The volumes of data generated by the OLTP server to the disk are comparatively much lower but this traffic requires faster response because transaction processing is a low latency application.

The point of congestion is the link between Switch 2 and the disk, for traffic from the switch to the disk. The return path is largely uncongested as there is little backup traffic on this path.

Service differentiation is needed at Switch 2 to prioritize the OLTP-server-to-disk traffic higher than the backup-server-to-disk traffic.

To configure traffic prioritization for the example application, follow these steps:

Step 1 Create the class maps.

```
Switch 2# config t
Switch 2(config)# qos class-map jc1 match-all
Switch 2(config-cmap)# match source-wwn 21:00:00:0c:50:02:ca:b5
Switch 2(config-cmap)# match destination-wwn 22:00:00:04:cf:22:eb:dc
Switch 2(config-cmap)# exit
Switch 2(config)# qos class-map jc2 match-all
Switch 2(config-cmap)# match source-wwn 21:00:00:0c:50:02:c7:ff
Switch 2(config-cmap)# match destination-wwn 22:00:00:04:cf:22:eb:dc
Switch 2(config-cmap)# exit
Switch 2(config)#
```

Step 2 Create the policy map.

```
Switch 2(config)# qos policy-map jp1
Switch 2(config-pmap)# class jc1
Switch 2(config-pmap-c)# priority high
Switch 2(config-pmap-c)# exit
Switch 2(config-pmap)# class jc2
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
Switch 2(config-pmap-c)# priority low
Switch 2(config-pmap-c)# exit
Switch 2(config-pmap)# exit
Switch 2(config)#
```

Step 3 Assign the service policy.

```
Switch 2(config)# qos service policy jp1 vsan 1
```

Step 4 Assign the weights for the DWRR queues.

```
Switch 2(config)# qos dwrr-q high weight 50
Switch 2(config)# qos dwrr-q medium weight 30
Switch 2(config)# qos dwrr-q low weight 20
```

Step 5 Repeat [Step 1](#) through [Step 4](#) on Switch 1 to address forward path congestion at both switches.

Congestion could occur anywhere in the example configuration. To address congestion of the return path at both switches, you need to create two more class maps and include them in the policy map as follows:

Step 1 Create two more class maps.

```
Switch 2(config)# qos class-map jc3 match-all
Switch 2(config-cmap)# match source-wwn 22:00:00:04:cf:22:eb:dc
Switch 2(config-cmap)# match destination-wwn 21:00:00:0c:50:02:ca:b5
Switch 2(config-cmap)# exit
Switch 2(config)# qos class-map jc4 match-all
Switch 2(config-cmap)# match source-wwn 22:00:00:04:cf:22:eb:dc
Switch 2(config-cmap)# match destination-wwn 21:00:00:0c:50:02:c7:ff
Switch 2(config-cmap)# exit
Switch 2(config)#
```

Step 2 Assign the class maps to the policy map.

```
Switch 2(config)# qos policy-map jp1
Switch 2(config-pmap)# class jc3
Switch 2(config-pmap-c)# priority high
Switch 2(config-pmap-c)# exit
Switch 2(config-pmap)# class jc4
Switch 2(config-pmap-c)# priority low
Switch 2(config-pmap-c)# exit
Switch 2(config-pmap)# exit
Switch 2(config)#
```

Step 3 Repeat [Step 1](#) through [Step 2](#) on Switch 1 to address return path congestion at both switches.

Ingress Port Rate Limiting

A port rate limiting feature is available in Cisco MDS SAN-OS Release 1.3. This feature helps control the bandwidth for individual FC ports. Port rate limiting is also referred to as ingress rate limiting because it controls ingress traffic into a FC port. The feature controls traffic flow by limiting the number of frames that are transmitted out of the exit point on the MAC. Port rate limiting works on all Fibre Channel ports. The rate limit ranges from 1 to 100% and the default is 100%.

Send documentation comments to mdsfeedback-doc@cisco.com.



Note Port rate limiting can only be configured on Cisco MDS 9100 Series switches, Cisco MDS 9216i switches, and MPS-14/2 modules.

This feature can only be configured if the QoS feature is enabled and if this configuration is performed on a Cisco MDS 9100 series switch, Cisco MDS 9216i switch, or MPS-14/2 module.

To configure the port rate limiting value, follow these steps.

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch # config t switch(config) # | Enters the configuration mode. |
| Step 2 | switch(config) # interface fc 1/1 | Selects the interface to specify the ingress port rate limit. |
| Step 3 | switch(config-if) # switchport ingress-rate 50 | Configures a 50% port rate limit for the selected interface. |
| | switch(config-if) # no switchport ingress-rate 50 | Reverts a previously configured rate to the factory default of 100%. |

Default Settings

Table 32-2 lists the default settings for FCC, QoS, and rate limiting features:

Table 32-2 Default FCC, QoS, and Rate Limiting Settings

| Parameters | Default |
|-------------------------|-----------|
| FCC protocol | Disabled. |
| QoS control traffic | Enabled. |
| QoS data traffic | Disabled. |
| Zone-based QoS priority | Low. |
| Rate limit | 100% |



Tracking and Redirecting Traffic

The Port Tracking feature is unique to the Cisco MDS 9000 Family of switches. This feature uses information about the operational state of the link to initiate a failure in the link that connects the edge device. This process of converting the indirect failure to a direct failure triggers a faster recovery process towards redundant links. When enabled, the port tracking feature brings down the configured links based on the failed link and forces the traffic to be redirected to another redundant link.

This chapter includes the following sections:

- [About Port Tracking, page 33-2](#)
- [Port Tracking Terminology, page 33-2](#)
- [Port Tracking Guidelines, page 33-3](#)
- [Port Tracking Features, page 33-3](#)
- [Enabling Port Tracking, page 33-3](#)
- [Configuring Linked Ports, page 33-3](#)
- [Displaying Port Tracking Information, page 33-6](#)
- [Default Settings, page 33-8](#)


Send documentation comments to mdsfeedback-doc@cisco.com.

About Port Tracking

Generally, hosts can instantly recover from a link failure on a link that is immediately (direct link) connected to a switch. However, recovering from an indirect link failure between switches in a WAN or MAN fabric with a keep-alive mechanism is dependent on several factors such as the time out values (TOVs) and on registered state change notification (RSCN) information (see the “[Fibre Channel Time Out Values](#)” section on page 39-2 and “[About RSCN Information](#)” section on page 18-7).

In [Figure 33-1](#), when the direct link 1 to the host fails, recovery can be immediate. However, when the ISL 2 fails between the two switches, recovery depends on TOVs, RSCNs, and other factors.

Figure 33-1 Traffic Recovery Using Port Tracking



As of Cisco SAN-OS Release 2.0(1b), the port tracking feature monitors and detects failures that cause topology changes and brings down the links connecting the attached devices. When you enable this feature and explicitly configure the linked and tracked ports, the Cisco SAN-OS software monitors the tracked ports and alters the operational state of the linked ports on detecting a link state change.

Port Tracking Terminology

The following terms are used in this chapter.

- **Tracked ports**—A port whose operational state is continuously monitored. The operational state of the tracked port is used to alter the operational state of one or more ports. Fibre Channel, VSAN, PortChannel, FCIP, or a Gigabit Ethernet port can be tracked. Generally, ports in E and TE port modes can also be Fx ports
- **Linked ports**—A port whose operational state is altered based on the operational state of the tracked ports. Only a Fibre Channel port can be linked.

Send documentation comments to mdsfeedback-doc@cisco.com.

Port Tracking Guidelines

Before configuring port tracking, consider the following guidelines:

- Verify that the tracked ports and the linked ports are on the same Cisco MDS switch.
- Be aware that the linked port is automatically brought down when the tracked port goes down.
- Do not track a linked port back to itself (for example, Port fc1/2 to Port fc2/5 and back to Port fc1/2) to avoid recursive dependency.

Port Tracking Features

Port tracking has the following features:

- The application brings the linked port down when the tracked port goes down. When the tracked port recovers from the failure and comes back up again, the tracked port is also brought up automatically (unless otherwise configured).
- You can forcefully continue to keep the linked port down, even though the tracked port comes back up. In this case, you must explicitly bring the port up when required.

Enabling Port Tracking

The port tracking feature is disabled by default in all switches in the Cisco 9000 Family. When you enable this feature, port tracking is globally enabled for the entire switch.

To configure port tracking, enable the port tracking feature and configure the linked port(s) for the tracked port.

To enable port tracking, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# port-track enable | Enables port tracking. |
| | switch(config)# no port-track enable | Removes the currently applied port tracking configuration and disables port tracking. |

Configuring Linked Ports

You can link ports using one of two methods:

- Operationally binding the linked port(s) to the tracked port (default)
- Continuing to keep the linked port down forcefully—even if the tracked port has recovered from the link failure.

Send documentation comments to mdsfeedback-doc@cisco.com.

Operational Binding

When you configure the first tracked port, operational binding is automatically in effect. When you use this method, you have the option to monitor multiple ports or monitor ports in one VSAN.

To operationally bind a tracked port, follow these steps:


| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc8/6 switch(config-if)# | Configures the specified interface and enters the interface configuration submenu. You can now configure tracked ports. Note This link symbolizes the direct link (1) in Figure 33-1 . |
| Step 3 | switch(config-if)# port-track interface port-channel 1 | Tracks interface fc8/6 with interface port-channel 1. When port-channel 1 goes down, interface fc8/6 is also brought down Note This link symbolizes the ISL (2) in Figure 33-1 . |
| | switch(config-if)# no port-track interface port-channel 1 | Removes the port tracking configuration that is currently applied to interface fc8/6. |

Tracking Multiple Ports

You can control the operational state of the linked port based on the operational states of multiple tracked ports. When more than one tracked port is associated with a linked port, the operational state of the linked port will be set to down only if all the associated tracked ports are down. Even if one tracked port is up, the linked port will stay up.

In [Figure 33-2](#), only if both ISLs 2 and 3 fail, will the direct link 1 be brought down. Direct link 1 will not be brought down if either 2 or 3 are still functioning as desired.

Figure 33-2 Traffic Recovery Using Port Tracking



Send documentation comments to mdsfeedback-doc@cisco.com.

To track multiple ports, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc8/6 | Configures the specified interface and enters the interface configuration submenu. You can now configure tracked ports. Note This link symbolizes the direct link (1) in Figure 33-2 . |
| Step 3 | switch(config-if)# port-track interface port-channel 1 | Tracks interface fc8/6 with interface port-channel 1. When port-channel 1 goes down, interface fc8/6 is also brought down. Note This link symbolizes the ISL (2) in Figure 33-2 . |
| Step 4 | switch(config-if)# port-track interface fcip 5 | Tracks interface fc8/6 with interface fcip 5. When FCIP 5 goes down, interface fc8/6 is also brought down. Note This link symbolizes the ISL (3) in Figure 33-2 . |

Monitoring Ports in a VSAN

You can optionally configure one VSAN from the set of all operational VSANs on the tracked port with the linked port by specifying the required VSAN. This level of flexibility provides higher granularity in tracked ports. In some cases, when a tracked port is a TE port, the set of operational VSANs on the port can change dynamically without bringing down the operational state of the port. In such cases, the port VSAN of the linked port can be monitored on the set of operational VSANs on the tracked port.

If you configure this feature, the linked port is up only when the VSAN is up on the tracked port.



Tip

The specified VSAN does not have to be the same as the port VSAN of the linked port.

To monitor a tracked port in a specific VSAN, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc8/6 | Configures the specified interface and enters the interface configuration submenu. You can now configure tracked ports. |
| Step 3 | switch(config-if)# port-track interface port-channel 1 vsan 2 | Enables tracking of the PortChannel in VSAN 2. |
| | switch(config-if)# no port-track interface port-channel 1 vsan 2 | Removes the VSAN association for the linked port. The PortChannel link remains in effect. |

Forceful Shutdown

If a tracked port flaps frequently, then tracking ports using the operational binding feature may cause frequent topology change. In this case, you may choose to keep the port in the down state until you are able to resolve the reason for these frequent flaps. Keeping the flapping port in the down state forces the traffic to flow through the redundant path until the primary tracked port problems are resolved. When the problems are resolved and the tracked port is back up, you can explicitly enable the interface.

Send documentation comments to mdsfeedback-doc@cisco.com.



Tip If you configure this feature, the linked port continues to remain in the shutdown state even after the tracked port comes back up. You must explicitly remove the forced shut state (by administratively bringing up this interface) of the linked port once the tracked port is up and stable.

To forcefully shutdown a tracked port, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/5 | Configures the specified interface and enters the interface configuration submenu. You can now configure tracked ports. |
| Step 3 | switch(config-if)# port-track force-shut | Forcefully shuts down the tracked port. |
| | switch(config-if)# no port-track force-shut | Removes the port shutdown configuration for the tracked port. |

Displaying Port Tracking Information

The **show** commands display the current port tracking settings for the Cisco MDS switch (see Examples 33-1 to 33-4).

Example 33-1 *Displays the Linked and Tracked Port Configuration*

```
switch# show interface
...
fc8/6 is down (All tracked ports down)      <-----Linked port
  Hardware is Fibre Channel, FCOT is short wave laser
  Port WWN is 21:c6:00:05:30:00:37:1e
  Admin port mode is auto, trunk mode is on
  Port vsan is 1
  Receive data field Size is 2112
  Beacon is turned off
  Port tracked with interface port-channel 1 vsan 2 (trunking) <-----Tracked port
  Port tracked with interface fcip 5 <-----Tracked port
  5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    269946 frames input, 22335204 bytes
      0 discards, 0 errors
      0 CRC, 0 unknown class
      0 too long, 0 too short
    205007 frames output, 10250904 bytes
      0 discards, 0 errors
    0 input OLS, 0 LRR, 0 NOS, 0 loop inits
    2 output OLS, 2 LRR, 0 NOS, 1 loop inits
    0 receive B2B credit remaining
    0 transmit B2B credit remaining
...
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 33-2 Displays a Tracked Port Configuration for a Fibre Channel Interface

```
switch# show interface fc1/1
fc1/1 is down (Administratively down)
  Hardware is Fibre Channel, FCOT is short wave laser w/o OFC (SN)
  Port WWN is 20:01:00:05:30:00:0d:de
  Admin port mode is FX
  Port vsan is 1
  Receive data field Size is 2112
  Beacon is turned off
Port tracked with interface fc1/2 (down)
Port tracked with interface port-channel 1 vsan 2 (down)
Port tracked with interface fcip1 (down)
5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  1 frames input, 128 bytes
    0 discards, 0 errors
    0 CRC, 0 unknown class
    0 too long, 0 too short
  1 frames output, 128 bytes
    0 discards, 0 errors
  0 input OLS, 0 LRR, 0 NOS, 0 loop inits
  0 output OLS, 0 LRR, 0 NOS, 0 loop inits
  0 receive B2B credit remaining
  0 transmit B2B credit remaining
```

Example 33-3 Displays a Tracked Port Configuration for a PortChannel Interface

```
switch# show interface port-channel 1
port-channel 1 is down (No operational members)
  Hardware is Fibre Channel
  Port WWN is 24:01:00:05:30:00:0d:de
  Admin port mode is auto, trunk mode is on
  Port vsan is 2
  Linked to 1 port(s)
Port linked to interface fc1/1
5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  0 frames input, 0 bytes
    0 discards, 0 errors
    0 CRC, 0 unknown class
    0 too long, 0 too short
  0 frames output, 0 bytes
    0 discards, 0 errors
  0 input OLS, 0 LRR, 0 NOS, 0 loop inits
  0 output OLS, 0 LRR, 0 NOS, 0 loop inits
No members
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 33-4 Displays a Forced Shut Configuration

```
switch# show interface fc 1/5

fc1/5 is up
  Hardware is Fibre Channel, FCOT is short wave laser
  Port WWN is 20:05:00:05:30:00:47:9e
  Admin port mode is F
  Port mode is F, FCID is 0x710005
  Port vsan is 1
  Speed is 1 Gbps
  Transmit B2B Credit is 64
  Receive B2B Credit is 16
  Receive data field Size is 2112
  Beacon is turned off
  Port track mode is force_shut <--this port remains shut even if the tracked port is back up
```

Default Settings

Table 33-1 lists the default settings for port tracking parameters.

Table 33-1 Default Port Tracking Parameters

| Parameters | Default |
|---------------------|----------------------------------|
| Port tracking | Disabled |
| Operational binding | Enabled along with port tracking |



Configuring the SAN Extension Tuner

The SAN extension tuner (SET) feature is unique to the Cisco MDS 9000 Family of switches. This feature helps you optimize FCIP performance by generating SCSI I/O commands and directing such traffic to a specific virtual target. You can specify the size of the test I/O transfers and how many concurrent I/Os to generate while testing. The SET reports the resulting I/Os per second (IOPS) and I/O latency, which helps you determine the number of concurrent I/Os needed to maximize FCIP throughput.

This chapter includes the following sections:

- [About SET, page 34-2](#)
- [License Prerequisites, page 34-2](#)
- [Tuner Guidelines, page 34-3](#)
- [Tuner Initialization, page 34-3](#)
- [Tuner Configuration, page 34-3](#)
- [nWWN Configuration, page 34-4](#)
- [Virtual N Port Configuration, page 34-5](#)
- [SCSI Read/Write Assignment, page 34-5](#)
- [Data Pattern, page 34-7](#)
- [Tuning Configuration Verification, page 34-8](#)
- [Default Settings, page 34-9](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

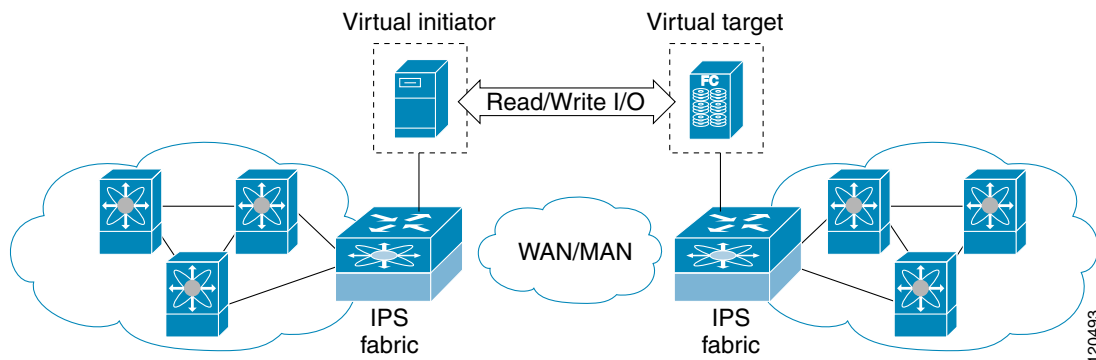
About SET

Applications such as remote copy and data backup use FCIP over an IP network to connect across geographically distributed SANs. To achieve maximum throughput performance across the fabric, you can tune the following configuration parameters:

- The TCP parameters for the FCIP profile, which includes the **max-bandwidth** parameter, the **min-available-bandwidth** parameter, and the **round-trip-time** (see the “[Window Management](#)” section on page 28-27).
- The number of concurrent SCSI I/Os generated by the application.
- The transfer size used by the application over an FCIP link.

As of Cisco SAN-OS Release 2.0(1b), SET is implemented in IPS ports. When enabled, this feature can be used to generate SCSI I/O commands (read and write) to the virtual target based on your configured options (see [Figure 34-1](#)).

Figure 34-1 SCSI Command Generation to the Virtual Target



The SET feature assist with tuning by generating varying SCSI traffic workloads. It also measures throughput and response time per I/ O over FCIP link.

License Prerequisites

To use the SET, you need to obtain the SAN_EXTN_OVER_IP license (see [Chapter 3, “Obtaining and Installing Licenses”](#)).

Send documentation comments to mdsfeedback-doc@cisco.com.

Tuner Guidelines

Before tuning the SAN fabric, be aware of the following guidelines:

- Be aware of the following implementation details:
 - The tuned configuration is not persistent.
 - The virtual N ports created do not register FC4 features supported with the name server. This is to avoid the hosts in the SAN from discovering these N ports as regular initiators or targets.
 - Login requests from other initiators in the SAN are rejected.
 - The virtual N ports do not implement the entire SCSI suite, it only implements the SCSI read and write commands.
 - Tuner initiators can only communicate with tuner targets.
- Verify that the Gigabit Ethernet interface is up at the physical layer (GBIC and Cable connected—an IP address is not required).
- Enable the iSCSI interface (no other iSCSI configuration is required)
- Configure the virtual N ports in a separate VSAN or zone as required by your network.
- Be aware that a separate VSAN with only virtual N ports is not required, but is recommended as some legacy HBAs may fail if logins to targets are rejected.
- Do not use same Gigabit Ethernet interface to configure virtual N ports and FCIP links—use different Gigabit Ethernet interfaces. While this is not a requirement, it is recommended as the traffic generated by the virtual N ports may interfere with the performance of the FCIP link.

Tuner Initialization

The tuning feature is disabled by default in all switches in the Cisco 9000 Family. When you enable this feature, tuning is globally enabled for the entire switch.

To enable the tuning feature, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# san-ext-tuner enable | Enables tuning. |
| | switch(config)# no san-ext-tuner enable | Removes the currently applied tuning configuration and disables tuning (default). |

Tuner Configuration

Figure 34-2 provides a sample physical setup in which the virtual N ports are created on ports that are not a part of the FCIP link for which the throughput and latency is measured.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 34-2 N Port Tuning Configuration Physical Example





Figure 34-2 provides a sample logical setup in which the virtual N ports are created on ports that are not a part of the FCIP link for which the throughput and latency is measured.

Figure 34-3 Logical Example of N Port Tuning for a FCIP Link



To tune the required FCIP link, follow these steps:

- Step 1

Configure the nWWN for the virtual N ports on the switch.
- Step 2

Configure the virtual N port on either side of the FCIP link.
- Step 3

Ensure that the virtual N ports are not visible to real initiators in the SAN. You can use zoning (see [Chapter 15, “Configuring and Managing Zones”](#)) or VSANs (see [Chapter 10, “Configuring and Managing VSANs”](#)) to segregate the real initiators.
- Step 4

Start the SCSI read and write I/Os.
- Step 5

Add more N ports (as required) to other Gigabit Ethernet ports in the switch to obtain maximum throughput. One scenario that may require additional N ports is if you use FCIP PortChannels.

nWWN Configuration

To configure the nWWNs for the tuner in this switch, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# san-ext-tuner switch(san-ext) # | Enters the SET configuration submode. |
| Step 2 | switch(san-ext) # nWWN 10:00:00:00:00:00:00:00 | Configures the nWWN for the SAN extension tuner. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Virtual N Port Configuration

To configure the virtual N port for tuning, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# san-ext-tuner switch(san-ext)# | Enters the SET configuration submenu. |
| Step 2 | switch(san-ext)# nWWN 10:00:00:00:00:00:00:00 | Configures the nWWN for the SAN extension tuner. |
| Step 3 | switch(san-ext)# nport pWWN 12:00:00:00:00:00:00:56 vsan 200 interface gigabitethernet 3/4 switch(san-ext-nport)# | Creates a virtual N port on the specified Gigabit Ethernet port and VSAN. This N port can act as a initiator or a target. |
| | switch(san-ext)# no nport pWWN 22:34:56:78:90:12:34:56 vsan 200 interface gigabitethernet 3/4 | Removes a virtual N port on the specified Gigabit Ethernet port and VSAN. |

SCSI Read/Write Assignment

You can assign SCSI read and write commands on a one-time basis or on a continuous basis.

To assign SCSI read and (or) write commands one a one-time basis, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# san-ext-tuner switch(san-ext)# | Enters the SET configuration submenu. |
| Step 2 | switch(san-ext)# nWWN 10:00:00:00:00:00:00:00 | Configures the nWWN for the SAN extension tuner. |
| Step 3 | switch(san-ext)# nport pWWN 12:00:00:00:00:00:00:56 vsan 200 interface gigabitethernet 3/4 switch(san-ext-nport)# | Creates a virtual N port on the specified Gigabit Ethernet port and VSAN. This N port can act as a initiator or a target. |
| Step 4 | switch(san-ext-nport)# read command-id 100 target 22:22:22:22:22:22:22:22 transfer-size 512000 outstanding-ios 2 num-transactions 5000000 | Specifies a transfer size of 512,000 bytes with two outstanding I/Os in the read command. The total number of I/Os is 5,000,000 bytes. |
| Step 5 | switch(san-ext-nport)# write command-id 100 target 22:22:22:22:22:22:22:22 transfer-size 512000 outstanding-ios 2 num-transactions 5000000 | Specifies a transfer size of 512,000 bytes with two outstanding I/Os in the write command received by the target. The total number of I/Os is 5,000,000 bytes. |
| Step 6 | switch(san-ext-nport)# stop command-id 100 switch(san-ext-nport)# stop all | Stops the command with the specified ID. Stops all outstanding commands. |
| Step 7 | switch(san-ext-nport)# clear counters | Clears the counters associated with this N port. |
| Step 8 | switch(san-ext-nport)# end switch# | Exits the SAN extension tuner submenu. |

Send documentation comments to mdsfeedback-doc@cisco.com.

To generate SCSI read or write commands continuously, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# san-ext-tuner switch(san-ext)# | Enters the SET configuration submode. |
| Step 2 | switch(san-ext)# nWWN 10:00:00:00:00:00:00:00 | Configures the nWWN for the SAN extension tuner. |
| Step 3 | switch(san-ext)# nport pWWN 12:00:00:00:00:00:00:56 vsan 200 interface gigabitethernet 3/4 switch(san-ext-nport)# | Creates a virtual N port on the specified Gigabit Ethernet port and VSAN. This N port can act as a initiator or a target. |
| Step 4 | switch(san-ext-nport)# read command-id 100 target 22:22:22:22:22:22:22:22 transfer-size 512000 outstanding-ios 2 continuous | Configures SCSI commands to be read continuously. Tip Use the stop command-id command to stop the outstanding configuration. |
| Step 5 | switch(san-ext-nport)# write command-id 100 target 22:22:22:22:22:22:22:22 transfer-size 512000 outstanding-ios 2 continuous | Configures SCSI commands to be written continuously. |
| Step 6 | switch(san-ext-nport)# stop command-id 100 switch(san-ext-nport)# stop command-id all | Stops the command with the specified ID. Stops all outstanding commands. |
| Step 7 | switch(san-ext-nport)# clear counters | Clears the counters associated with this N port. |
| Step 8 | switch(san-ext-nport)# end switch# | Exits the SAN extension tuner submode. |

To specify a transfer ready size for a SCSI write command, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# san-ext-tuner switch(san-ext)# | Enters the SET configuration submode. |
| Step 2 | switch(san-ext)# nWWN 10:00:00:00:00:00:00:00 | Configures the nWWN for the SAN extension tuner. |
| Step 3 | switch(san-ext)# nport pWWN 12:00:00:00:00:00:00:56 vsan 200 interface gigabitethernet 3/4 switch(san-ext-nport)# | Creates a virtual N port on the specified Gigabit Ethernet port and VSAN. This N port can act as a initiator or a target. |
| Step 4 | switch(san-ext-nport)# write command-id 100 target 22:22:22:22:22:22:22:22 transfer-size 512000 outstanding-ios 2 num-transactions 5000000 | Specifies a transfer size of 512,000 bytes with two outstanding I/Os in the write command received by the target. The total number of I/Os is 5,000,000 bytes. |
| Step 5 | switch(san-ext-nport)# transfer-ready-size 512000 | Specifies the maximum transfer ready size of 512,000 bytes as a target for SCSI write commands. For a SCSI write command with a larger size, the target performs multiple transfers based on the specified transfer size. |
| | switch(san-ext-nport)# no transfer-ready-size 512000 | Removes the specified transfer ready size configuration for SCSI write commands. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|---|--|
| Step 6 | switch(san-ext-nport)# stop command-id 100 | Stops the command with the specified ID. |
| Step 7 | switch(san-ext-nport)# end switch# | Exits the SAN extension tuner submode. |

Data Pattern

By default, an all-zero pattern is used as the pattern for data generated by the virtual N ports. You can optionally specify a file as the data pattern to be generated by selecting a data pattern file from one of three locations: the bootflash: directory, the volatile: directory, or the slot0: directory. This option is especially useful when testing compression over FCIP links. You can also use Canterbury corpus or artificial corpus files for benchmarking purposes.

To optionally configure a data pattern for SCSI commands, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# san-ext-tuner switch(san-ext)# | Enters the SET configuration submode. |
| Step 2 | switch(san-ext)# nport pWWN 12:00:00:00:00:00:56 vsan 200 interface gigabitethernet 3/4 switch(san-ext-nport)# | Creates a virtual N port on the specified Gigabit Ethernet port and VSAN. This N port can act as a initiator or a target. |
| Step 3 | switch(san-ext-nport)# data-pattern-file bootflash://DataPatternFile | Specifies the data pattern used by the N port to generate data as a target for read commands and initiator for write commands. |
| | switch(san-ext-nport)# no data-pattern-file | Removes the specified transfer ready size configuration for SCSI write commands and defaults to using the all-zero pattern. |
| Step 4 | switch(san-ext-nport)# write command-id 100 target 22:22:22:22:22:22:22:22 transfer-size 512000 outstanding-ios 2 num-transactions 5000000 | Specifies a transfer size of 512,000 bytes with two outstanding I/Os. The total number of I/Os is 5,000,000 bytes. |
| Step 5 | switch(san-ext-nport)# stop command-id 100 | Stops the command with the specified ID. |
| Step 6 | switch(san-ext-nport)# clear counters | Clears the counters associated with this N port. |
| Step 7 | switch(san-ext-nport)# end switch# | Exits the SAN extension tuner submode. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Tuning Configuration Verification

The **show** commands display the current tuning settings for the Cisco MDS switch (see Examples 34-1 to 34-6).

Example 34-1 *Displays Entries in the FLOGI Database*

```
switch# show flogi database
-----
INTERFACE  VSAN    FCID      PORT NAME                               NODE NAME
-----
iscsi3/4   200      0x050000  12:00:00:00:00:00:00:56               10:00:00:00:00:00:00:00
```

Example 34-2 *Displays Details for a VSAN Entry in the FLOGI Database*

```
switch# show fcns database vsan 200
VSAN 200
-----
FCID      TYPE    PWWN      (VENDOR)      FC4-TYPE:FEATURE
-----
0x020000  N       22:22:22:22:22:22:22:22  scsi-fcp
0x050000  N       12:00:00:00:00:00:00:56  scsi-fcp
```

Example 34-3 *Displays All Virtual N Ports Configured on the Specified Interface*

```
switch# show san-ext-tuner interface gigabitethernet 3/4 nport pwwn
12:00:00:00:00:00:00:56 vsan 200 counters
Statistics for nport
Node name 10:00:00:00:00:00:00:00 Port name 12:00:00:00:00:00:00:56
I/Os per second      : 148
  Read                : 0%
  Write               : 100%
Ingress MB per second : 0.02 MBs/sec (Max -0.02 MBs/sec)
Egress MB per second  : 73.97 MBs/sec (Max -75.47 MBs/sec)
Average Response time per I/O : Read - 0 us, Write - 13432 us
Maximum Response time per I/O : Read - 0 us, Write - 6953 us
Minimum Response time per I/O : Read - 0 us, Write - 19752 us
Errors                : 0
```

Example 34-4 *Displays N ports Configured on a Specified Gigabit Ethernet Interface.*

```
switch# show san-ext-tuner interface gigabitethernet 3/1
-----
Interface      NODE NAME                               PORT NAME                               VSAN
-----
GigabitEthernet3/1  10:00:00:00:00:00:00:00  10:00:00:00:00:00:00:01  91
```

Example 34-5 *Displays the Transfer Ready Size Configured for a Specified N Port*

```
switch# show san-ext-tuner interface gigabitethernet 3/1 nport pwwn 10:0:0:0:0:0:1 vsan
91
Node name      : 10:00:00:00:00:00:00:00
Port name      : 10:00:00:00:00:00:00:01
Transfer ready size : all
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 34-6 Displays All Virtual N Ports Configured in This Switch

```
switch# show san-ext-tuner nports
```

```
-----
Interface          NODE NAME          PORT NAME          VSAN
-----
GigabitEthernet3/1  10:00:00:00:00:00:00  10:00:00:00:00:00:01  91
```

Default Settings

Table 34-1 lists the default settings for tuning parameters.

Table 34-1 Default Tuning Parameters

| Parameters | Default |
|------------------------|---|
| Tuning | Disabled. |
| Transfer ready size | Same as the transfer size in the SCSI write command. |
| Outstanding I/Os | 1. |
| Number of transactions | 1. |
| Data generation format | All-zero format. |

Send documentation comments to mdsfeedback-doc@cisco.com.



Scheduling Maintenance Jobs

The Cisco MDS command scheduler feature helps you schedule configuration and maintenance jobs in any switch in the Cisco MDS 9000 Family. This feature is available in the Cisco SAN-OS Release 2.0(1b) software. You can use this feature to schedule jobs on a one-time basis or periodically.

This chapter includes the following sections:

- [About the Command Scheduler, page 35-1](#)
- [Scheduler Terminology, page 35-1](#)
- [Scheduling Guidelines, page 35-2](#)
- [Scheduler Configuration, page 35-2](#)
- [Scheduler Configuration Verification, page 35-8](#)
- [Default Settings, page 35-9](#)

About the Command Scheduler

The MDS command scheduler provides a facility to schedule a job (set of CLI commands) or multiple jobs at a specified time in the future. The job(s) can be executed once at a specified time in the future or at periodic intervals.



Note

To use the command scheduler, you do not need to obtain any license. This feature is available in all switches in the Cisco MDS family that use the Cisco SAN-OS Release 2.0(1b) software.

You can use this feature to schedule zone set changes, QOS policy changes, backup data, save the configuration and other similar jobs.

Scheduler Terminology

The following terms are used in this chapter.

- **Job**—A job is a set of SAN-OS CLI commands (EXEC and config mode) that are executed as defined in the schedule.
- **Schedule**—A schedule determines the time when the assigned jobs must be executed. Multiple jobs can be assigned to a schedule. A schedule executes in one of two modes: one-time or periodic.

Send documentation comments to mdsfeedback-doc@cisco.com.

- Periodic mode—A job is executed at the user-specified periodic intervals, until it is deleted by the administrator. The following types of periodic intervals are supported:
 - Daily—The job is executed once a day
 - Weekly—The job is executed once a week
 - Monthly—The job is executed once a month
 - Delta—The job is executed beginning at the specified start time and thereafter at user-specified intervals (days:hours:minutes).
- One-time mode—The job is executed once at a user-specified time.

Scheduling Guidelines

Before scheduling jobs on a Cisco MDS switch, be aware of the following guidelines:

- A user who is authenticated and authorized by a remote service (for example, RADIUS) cannot schedule jobs.
- Be aware that the scheduled job can fail if it encounters one of the following situations when executing the job:
 - If the license has expired for a feature at the time when a job containing commands pertaining to that feature is scheduled.
 - If a feature is disabled at the time when a job containing commands pertaining to that feature is scheduled.
 - If you have removed a module from a slot and the job has commands pertaining to the interfaces for that module or slot.
- Verify that you have configured the time. The scheduler does not have any default time configured. If you create a schedule and assign job(s) and do not configure the time, that schedule is not launched.
- While defining a job, verify that no interactive or disruptive commands (for example, **copy bootflash: file ftp: URI, write erase**, and other similar commands) are specified as part of a job since the job is executed noninteractively at the scheduled time.

Scheduler Configuration

To configure the command scheduler, follow these steps:

-
- | | |
|---------------|--|
| Step 1 | Enable (initialize) the scheduler. |
| Step 2 | Define the job. |
| Step 3 | Define the schedule and assign jobs to the schedule. |
| Step 4 | Specify the time for the schedule(s). |
| Step 5 | Verify the scheduled configuration. |
-

Send documentation comments to mdsfeedback-doc@cisco.com.

Command Scheduler Initialization

To use the scheduling feature, you must explicitly enable this feature on the required switches in the fabric. By default, this feature is disabled in all switches in the Cisco MDS 9000 family.

The configuration and verification commands for the command scheduler feature are only available when this feature is enabled on a switch. When you disable this feature, all related configurations are automatically discarded.

To enable the command scheduling feature, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# scheduler enable | Enables the scheduler. |
| | switch(config)# no scheduler enable | Discards the scheduler configuration and disables the scheduler (default). |

Job Definition

To define a job, you must specify the job name. This action places you in the job definition (`config-job`) submode. In this submode, you can define the sequence of CLI commands that the job has to perform. Be sure to exit the `config-job` submode to complete the job definition.



Caution

You cannot modify or remove a command after entering the sequence of commands. To make changes, you must explicitly delete the defined job name and restart this process.



Note

You must exit the `config-job` submode for the job definition to be complete.

Send documentation comments to mdsfeedback-doc@cisco.com.

To define a job for the command scheduler, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# conf t switch(config)# | Enters the configuration mode. |
| Step 2 | switch(config)# scheduler job name addMemVsan99 switch(config-job)# | Defines a job name and enters the job definition submenu |
| Step 3 | switch(config-job)# config terminal switch(config-job-config)# vsan database switch(config-job-config-vsan-db)# vsan 99 interface fc1/1 - 4 switch(config-job-config-vsan-db)# end switch# | Specifies a sequence of actions for the specified job. The defined commands are checked for validity and stored for future use. Note Be sure you exit the config-job submenu. |
| | switch(config)# scheduler job name offpeakQOS switch(config-job)# conf t switch(config-job-config)# qos class-map offpeakbackupcmap match-all switch(config-job-config-cmap)# match source-wwn 23:15:00:05:30:00:2a:1f switch(config-job-config-cmap)# match destination-wwn 20:01:00:05:30:00:28:df switch(config-job-config-cmap)# exit switch(config-job-config)# qos policy-map offpeakbackuppolicy switch(config-job-config-pmap)# class offpeakbackupcmap switch(config-job-config-pmap-c)# priority high switch(config-job-config-pmap-c)# exit switch(config-job-config-pmap)# exit switch(config-job-config)# qos service policy offpeakbackuppolicy vsan 1 switch(config-job-config)# end switch# | Provides another example of scheduling a different set of jobs. |

Job Deletion

To delete a job for the command scheduler, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# conf t switch(config)# | Enters the configuration mode. |
| Step 2 | switch(config)# no scheduler job name addMemVsan99 | Deletes a defined job and all commands defined within that job. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Schedule Definition

After defining jobs, you can create schedules and assign jobs to the schedule. Subsequently, you can configure the time of execution. The execution can be one-time or periodic depending on your requirements. If the time for the schedule is not configured, then it will never be executed.

Periodic Schedule Definition

When you specify a periodic job execution, that job is executed periodically at the specified (daily, weekly, monthly, or delta) intervals.

To specify a periodic job for the command scheduler, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# conf t switch(config)# | Enters the configuration mode. |
| Step 2 | switch(config)# scheduler schedule name weekendbackup qos switch(config-schedule)# | Defines a job schedule (weekendbackup) and enters the submode for that schedule. |
| | switch(config)# no scheduler schedule name weekendbackup | Deletes the defined schedule. |
| Step 3 | switch(config-schedule)# job name offpeakZoning switch(config-schedule)# job name offpeakQOS | Assign two jobs offpeakZoning and offpeakQOS for this schedule. |
| | switch(config-schedule)# no job name addMem99 | Deletes the job assigned for this schedule. |
| Note | The following examples are only provided for reference. | |
| Step 4 | switch(config-schedule)# time daily 23:00 | Executes the specified jobs at 11 p.m. every day. |
| | switch(config-schedule)# time weekly Sun:23:00 | Specifies a weekly execution every Sunday at 11 p.m. |
| | switch(config-schedule)# time monthly 28:23:00 | Specifies a monthly execution at 11 p.m. on the 28th of each month. If you specify the date as either 29, 30, or 31, the command is automatically executed on the last day of each month. |
| | switch(config-schedule)# time start now repeat 48:00 | Specifies a job to be executed every 48 hours beginning 2 minutes from <i>now</i> —if today is September 24th 2004 and the time is now 2:00 p.m., the command begins executing at 2 minutes past 2:00 p.m. on September 24th 2004 and continues to execute every 48 hours after that. |
| | switch(config-schedule)# time start 14:00 repeat 14:00:00 | If today is September 24th, 2004 (Friday), this command specifies the job to be executed every alternate Friday at 2 p.m. (every 14 days). |

Send documentation comments to mdsfeedback-doc@cisco.com.

The most significant fields in the **time** parameter are optional. If you omit the most significant fields, the values are assumed to be the same as the current time. For example, if the current time is September 24th, 2004, 22:00 hours, then the commands are executed as follows:

- The **time start 23:00 repeat 4:00:00** command implies a start time of September 24th, 2004 23:00 hours.
- The **time daily 55** command implies every day at 22:55 hours.
- The **time weekly 23:00** command implies every Friday at 23:00 hours.
- The **time monthly 23:00** command implies the 24th of every month at 23:00 hours.



Note

If the time interval configured for any schedule is smaller than the time taken to execute its assigned job(s), then the subsequent schedule execution occurs only after the configured interval amount of time has elapsed following the completion time of the last iteration of the schedule. For example, a schedule is executed at 1-minute intervals and a job assigned to it takes 2 minutes to complete. If the first schedule is at 22:00 hours, the job finishes at 22:02 following which, the 1-minute interval is observed and the next execution occurs at 22:03 which finishes at 22:05.

One-Time Schedule Definition

When you specify a one-time job execution, that job is only executed once

To specify a one-time job for the command scheduler, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# conf t switch(config)# | Enters the configuration mode. |
| Step 2 | switch(config)# scheduler schedule name configureVsan99 switch(config-schedule)# | Defines a job schedule (configureVsan99) and enters the submode for that schedule. |
| Step 3 | switch(config-schedule)# job name addMemVsan99 | Assigns a predefined job name (addMemVsan99) for this schedule. |
| Step 4 | switch(config-schedule)# time start 2004:12:14:23:00 | Specifies a one-time execution on December 14th, 2004 at 11 p.m. |
| | switch(config-schedule)# no time | Deletes the time assigned for this schedule. |

Schedule Deletion

To delete a schedule, follow these steps:

| | Command | Purpose |
|--------|---|--------------------------------|
| Step 1 | switch# conf t switch(config)# | Enters the configuration mode. |
| Step 2 | switch(config)# no scheduler schedule name weekendbackup | Deletes the defined schedule. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Job Disassociation

To disassociate an assigned job, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# conf t switch(config)# | Enters the configuration mode. |
| Step 2 | switch(config)# scheduler schedule name weekendbackupqos switch(config-schedule)# | Defines a job schedule (weekendbackup) and enters the submode for that schedule. |
| Step 3 | switch(config-schedule)# no job name addMem99 | Deletes the job assigned for this schedule. |

Schedule Time Deletion

To delete the schedule time, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# conf t switch(config)# | Enters the configuration mode. |
| Step 2 | switch(config)# scheduler schedule name weekendbackupqos switch(config-schedule)# | Defines a job schedule (weekendbackup) and enters the submode for that schedule. |
| Step 3 | switch(config-schedule)# no time | This will delete the schedule time configuration. The schedule will not be run until the time is configured again. |

Execution Log

The command scheduler maintains a log file. While you cannot modify the contents of this file, you can change the file size. This log file is a circular log which contains the output of the job executed. If the output of the job is greater than the log file, then the output stored in this file remains truncated.

You can configure the log file size to be a maximum of 1024KB. The default size of the execution log file is 16KB.

To configure the execution log file size, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# conf t switch(config)# | Enters the configuration mode. |
| Step 2 | switch(config)# scheduler logfile size 1024 | Configures the log file to be a maximum of 1024 KB |
| | switch(config)# no scheduler logfile size | Defaults to the log size of 16KB. |

Clearing the Log File Contents

To clear the contents of the scheduler log file, issue the **clear scheduler logfile** command in EXEC mode.

```
switch# clear scheduler logfile
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Scheduler Configuration Verification

The **show** commands display the current command scheduler settings for the Cisco MDS switch (see Examples 35-1 to 35-4).

Example 35-1 Displays the Commands to be Executed for a Specified Job

```
switch# show scheduler job addMemVsan99
Job Name: addMemVsan99
-----
config terminal
vsan database
  vsan 99 interface fc1/1
  vsan 99 interface fc1/2
  vsan 99 interface fc1/3
  vsan 99 interface fc1/4
```

Example 35-2 Displays the Execution Status of the Schedule

```
switch# show scheduler schedule configureVsan99
Schedule Name      : configureVsan99
-----
User Name          : admin
Schedule Type      : Run once on Tue Aug 10 09:48:00 2004
Last Execution Time: Tue Aug 10 09:48:00 2004
-----
      Job Name                      Status
-----
addMemVsan99                      Success (0)
```

Example 35-3 Displays the Execution Log of All Jobs Executed in the System

```
switch# show scheduler logfile
Job Name      : addMemVsan99          Job Status: Success (0)
Schedule Name : configureVsan99      User Name : admin
Completion time: Tue Aug 10 09:48:00 2004
----- Job Output -----
`config terminal`
`vsan database`
`vsan 99 interface fc1/1`
`vsan 99 interface fc1/2`
`vsan 99 interface fc1/3`
`vsan 99 interface fc1/4`
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 35-4 *Displays the Scheduler Configuration on the Switch*

```
switch# show scheduler config
config terminal
  scheduler enable
  scheduler logfile size 512
end

config terminal
  scheduler job name addMemVsan99
  config terminal
    vsan database
    vsan 99 interface fc1/1
    vsan 99 interface fc1/2
    vsan 99 interface fc1/3
    vsan 99 interface fc1/4
  end

config terminal
  scheduler schedule name configureVsan99
    time start 2004:8:10:9:52
    job name addMemVsan99
  end
```

Default Settings

Table 35-1 lists the default settings for command scheduling parameters.

Table 35-1 *Default Command Scheduler Parameters*

| Parameters | Default |
|-------------------|-----------|
| Command scheduler | Disabled. |
| Log file size | 16 KB |

Send documentation comments to mdsfeedback-doc@cisco.com.



Configuring System Message Logging

This chapter describes how to configure system message logging on Cisco MDS 9000 Family switches. It includes the following sections:

- [About System Message Logging, page 36-1](#)
- [System Message Logging Configuration, page 36-3](#)
- [System Message Logging Configuration Distribution, page 36-7](#)
- [Displaying System Message Logging Information, page 36-9](#)
- [Default Settings, page 36-13](#)

About System Message Logging

The system message logging software saves messages in a log file or directs the messages to other devices. This feature provides you with the following capabilities:

- Provides logging information for monitoring and troubleshooting
- Allows you to select the types of captured logging information.
- Allows you to select the destination server to forward the captured logging information.

By default, the switch logs normal but significant system messages to a log file and sends these messages to the system console. You can specify which system messages should be saved based on the type of facility (see [Table 36-1](#)) and the severity level (see [Table 36-2](#)). Messages are time-stamped to enhance real-time debugging and management.

You can access logged system messages using the CLI or by saving them to a properly configured system message logging server. The switch software saves system messages in a file that can be configured to save up to 4 MB. You can monitor system messages remotely by accessing the switch through Telnet, SSH, or the console port, or by viewing the logs on a system message logging server.



Note

When the switch first initializes, the network is not connected until initialization completes. Therefore, messages are not redirected to a system message logging server for a few seconds.

Log messages are not saved across system reboots. However, a maximum of 100 log messages with a severity level of critical and below (levels 0, 1, and 2) are saved in NVRAM.

[Table 36-1](#) describes some samples of the facilities supported by the system message logs.

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 36-1 Internal Logging Facilities

| Facility Keyword | Description | Standard or Cisco MDS Specific |
|-------------------------|--------------------------------|--------------------------------|
| acl | ACL manager | Cisco MDS 9000 Family specific |
| all | All facilities | Cisco MDS 9000 Family specific |
| auth | Authorization system | Standard |
| authpriv | Authorization (private) system | Standard |
| bootvar | Bootvar | Cisco MDS 9000 Family specific |
| callhome | Call Home | Cisco MDS 9000 Family specific |
| cron | Cron or at facility | Standard |
| daemon | System daemons | Standard |
| fcc | FCC | Cisco MDS 9000 Family specific |
| fcdomain | fcdomain | Cisco MDS 9000 Family specific |
| fcns | Name server | Cisco MDS 9000 Family specific |
| fcs | FCS | Cisco MDS 9000 Family specific |
| flogi | FLOGI | Cisco MDS 9000 Family specific |
| fspf | FSPF | Cisco MDS 9000 Family specific |
| ftp | File Transfer Protocol | Standard |
| ipconf | IP configuration | Cisco MDS 9000 Family specific |
| ipfc | IPFC | Cisco MDS 9000 Family specific |
| kernel | Kernel | Standard |
| local0 to local7 | Locally defined messages | Standard |
| lpr | Line printer system | Standard |
| mail | Mail system | Standard |
| mcast | Multicast | Cisco MDS 9000 Family specific |
| module | Switching module | Cisco MDS 9000 Family specific |
| news | USENET news | Standard |
| ntp | NTP | Cisco MDS 9000 Family specific |
| platform | Platform manager | Cisco MDS 9000 Family specific |
| port | Port | Cisco MDS 9000 Family specific |
| port-channel | PortChannel | Cisco MDS 9000 Family specific |
| qos | QoS | Cisco MDS 9000 Family specific |
| rdl | RDL | Cisco MDS 9000 Family specific |
| rib | RIB | Cisco MDS 9000 Family specific |
| rscn | RSCN | Cisco MDS 9000 Family specific |
| securityd | Security | Cisco MDS 9000 Family specific |
| syslog | Internal system messages | Standard |
| sysmgr | System manager | Cisco MDS 9000 Family specific |
| tlport | TL port | Cisco MDS 9000 Family specific |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 36-1 ***Internal Logging Facilities (continued)***

| Facility Keyword | Description | Standard or Cisco MDS Specific |
|------------------|----------------------------------|--------------------------------|
| user | User process | Standard |
| uucp | UNIX-to-UNIX Copy Program | Standard |
| vhbad | Virtual host base adapter daemon | Cisco MDS 9000 Family specific |
| vni | Virtual network interface | Cisco MDS 9000 Family specific |
| vrrp_cfg | VRRP configuration | Cisco MDS 9000 Family specific |
| vrrp_eng | VRRP engine | Cisco MDS 9000 Family specific |
| vsan | VSAN system messages | Cisco MDS 9000 Family specific |
| vshd | vshd | Cisco MDS 9000 Family specific |
| wwn | WWN manager | Cisco MDS 9000 Family specific |
| xbar | Xbar system messages | Cisco MDS 9000 Family specific |
| zone | Zone server | Cisco MDS 9000 Family specific |

Table 36-2 describes the severity levels supported by the system message logs.

Table 36-2 ***Error Message Severity Levels***

| Level Keyword | Level | Description | System Message Definition |
|----------------------|-------|----------------------------------|---------------------------|
| emergencies | 0 | System unusable | LOG_EMERG |
| alerts | 1 | Immediate action needed | LOG_ALERT |
| critical | 2 | Critical conditions | LOG_CRIT |
| errors | 3 | Error conditions | LOG_ERR |
| warnings | 4 | Warning conditions | LOG_WARNING |
| notifications | 5 | Normal but significant condition | LOG_NOTICE |
| informational | 6 | Informational messages only | LOG_INFO |
| debugging | 7 | Debugging messages | LOG_DEBUG |



Note

Refer to the *Cisco MDS 9000 Family System Messages References* for details on the error log message format.

System Message Logging Configuration

System logging messages are sent to the console based on the default (or configured) logging facility and severity values.

Send documentation comments to mdsfeedback-doc@cisco.com.

Message Logging Initiation

You can disable logging to the console or enable logging to a given Telnet or SSH session.

- When you disable or enable logging to a console session, that state is applied to all future console sessions. If you exit and log in again to a new session, the state is preserved.
- When you enable or disable logging to a Telnet or SSH session, that state is applied only to that session. If you exit and log in again to a new session, the state is not preserved.

To enable or disable the logging state for a Telnet, or SSH session, follow these steps:

| | Command | Purpose |
|--------|------------------------------------|---|
| Step 1 | switch# terminal monitor | Enables logging for a Telnet, or SSH session. Note A console session is enabled by default. |
| Step 2 | switch# terminal no monitor | Disables logging for a Telnet, or SSH session. Note A Telnet or SSH session is disabled by default. |

Console Severity Level

When logging is enabled for a console session (default), you can configure the severity levels of messages that appear on the console. The default severity for console logging is 2 (critical).



Tip

The current critical (default) logging level is maintained if the console baud speed is 9600 baud (default). All attempts to change the console logging level generates an error message. To increase the logging level (above critical), you must change the console baud speed to 38400 baud (see the [“Configuring Console Port Settings”](#) section on page 4-34).

To configure the severity level for a logging facility, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# logging console 3 | Configures console logging at level 3 (error). Logging messages with a severity level of 3 or above are displayed on the console. |
| | switch(config)# logging console | Reverts console logging to the factory set default severity level of 2 (critical). Logging messages with a severity level of 2 or above are displayed on the console. |

Module Logging

By default, logging is enabled at level 7 for all modules. You can enable or disable logging for each module at a specified level.

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure the severity level for a logging facility, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# logging module 1 | Configures module logging at level 1 (alerts). |
| | switch(config)# logging module | Configures module logging for all modules in the switch. |
| | switch(config)# no logging module | Reverts module logging to the factory set default of not configuring logging for all modules. |

Facility Severity Level

To configure the severity level for a logging facility, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# logging level kernel 4 | Configures Telnet or SSH logging for the kernel facility at level 4 (warning). As a result, logging messages with a severity level of 4 or above are displayed. |

Log Files

Logging messages may be saved to a log file. You can configure the name of this file and restrict its size as required. The default log file name is messages. The file name can have up to 80 characters and the file size ranges from 4096 bytes to 4194304 bytes.

You can rename this file using the **logging logfile** command.

To send log messages to file, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# logging logfile ManagerLog 3 size 3000000 | Configures log information for errors or events above severity level 3 to be logged in a file named ManagerLog. By configuring a size, you are restricting the file size to 3,000,000 bytes. |
| | switch(config)# logging logfile ManagerLog 3 | Turns off logging to a file named ManagerLog and reverts to the default messages file. |

The configured log file is saved in the /var/log/external directory. The location of the log file cannot be changed. You can use the **show logging logfile** and **clear logging logfile** commands to view and delete the contents of this file. As of Cisco MDS SAN-OS Release 2.1(1a), you can use the **dir log:** command to view logging file statistics. You can use the **delete log:** command to remove the log file.

You can copy the logfile to a different location using the **copy log:** command using additional copy syntax (see the “[Copying Files](#)” section on page 4-28).

Send documentation comments to mdsfeedback-doc@cisco.com.

System Message Logging Servers

You can configure a maximum of three system message logging servers.

To send log messages to a UNIX system message logging server, you must configure the system message logging daemon on a UNIX server. Log in as root, and follow these steps:

Step 1 Add the following line to the `/etc/syslog.conf` file.

```
local1.debug                /var/log/myfile.log
```



Note Be sure to add five tab characters between **local1.debug** and **/var/log/myfile.log**. Refer to entries in the `/etc/syslog.conf` file for further examples.

The switch sends messages according to the specified facility types and severity levels. The **local1** keyword specifies the UNIX logging facility used. The messages from the switch are generated by user processes. The **debug** keyword specifies the severity level of the condition being logged. You can set UNIX systems to receive all messages from the switch.

Step 2 Create the log file by entering these commands at the UNIX shell prompt:

```
$ touch /var/log/myfile.log
$ chmod 666 /var/log/myfile.log
```

Step 3 Make sure the system message logging daemon reads the new changes by entering this command:

```
$ kill -HUP ~cat /etc/syslog.pid~
```

To configure system message logging servers, follow these steps:

| | Command | Purpose |
|---------------|--|--|
| Step 1 | switch# config t switch# | Enters configuration mode. |
| Step 2 | switch(config)# logging server 172.22.00.00 | Configures the switch to forward log messages according to the specified facility types and severity levels to remote multiple servers specified by its hostname or IP address (172.22.00.00). |
| | switch(config)# logging server 172.22.00.00 facility local1 | Configures the switch to forward log messages according to the specified facility (local1) for the server IP address (172.22.00.00). The default outgoing facility is local7. |
| | switch(config)# no logging server 172.11.00.00 | Removes the specified server (172.11.00.00) and reverts to factory default. |

Outgoing System Message Logging Server Facilities

All system messages have a logging facility and a level. The logging facility can be thought of as *where* and the level can be thought of as *what*.

The single system message logging daemon (syslogd) sends the information based on the configured **facility** option. If no facility is specified, local7 is the default outgoing facility.

The internal facilities are listed in [Table 36-1](#) and the outgoing logging facilities are listed in [Table 36-3](#).

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 36-3 **Outgoing Logging Facilities**

| Facility Keyword | Description | Standard or Cisco MDS Specific |
|-------------------------|--------------------------------|----------------------------------|
| auth | Authorization system | Standard |
| authpriv | Authorization (private) system | Standard |
| cron | Cron or at facility | Standard |
| daemon | System daemons | Standard |
| ftp | File Transfer Protocol | Standard |
| kernel | Kernel | Standard |
| local0 to local7 | Locally defined messages | Standard (local7 is the default) |
| lpr | Line printer system | Standard |
| mail | Mail system | Standard |
| news | USENET news | Standard |
| syslog | Internal system messages | Standard |
| user | User process | Standard |
| uucp | UNIX-to-UNIX Copy Program | Standard |

System Message Logging Configuration Distribution

As of Cisco SAN-OS Release 2.0(1b), you can enable fabric distribution for all Cisco MDS switches in the fabric. When you perform system message logging configurations, and distribution is enabled, that configuration is distributed to all the switches in the fabric.

You automatically acquire a fabric-wide lock when you issue the first configuration command after you enabled distribution in a switch. The system message logging server uses the effective and pending database model to store or commit the commands based on your configuration. When you commit the configuration changes, the effective database is overwritten by the configuration changes in the pending database and all the switches in the fabric receive the same configuration. After making the configuration changes, you can choose to discard the changes by aborting the changes instead of committing them. In either case, the lock is released. Refer to [Chapter 9, “Using the CFS Infrastructure”](#) for more information on the CFS application.

To enable fabric distribution for system message logging server configurations, follow these steps:

| | Command | Purpose |
|---------------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# logging distribute | Enables the system message logging server configuration to be distributed to all switches in the fabric, acquires a lock, and stores all future configuration changes in the pending database. |
| | switch(config)# no logging distribute | Disables (default) system message logging server configuration distribution to all switches in the fabric. |

Send documentation comments to mdsfeedback-doc@cisco.com.

To commit the system message logging server configuration changes, follow these steps:

| | Command | Purpose |
|--------|---------------------------------------|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# logging commit | Distributes the configuration changes to all switches in the fabric, releases the lock, and overwrites the effective database with the changes made to the pending database. |

To discard the system message logging server configuration changes, follow these steps:

| | Command | Purpose |
|--------|--------------------------------------|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# logging abort | Discards the system message logging server configuration changes in the pending database and releases the fabric lock. |

Fabric Lock Override

If you have performed a system message logging task and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.



Tip

The changes are only available in the volatile directory and are subject to being discarded if the switch is restarted.

To use administrative privileges and release a locked system message logging session, use the **clear logging session** command.

```
switch# clear logging session
```

Database Merge Guidelines

See the “CFS Merge Support” section on page 9-7 for detailed concepts.

When merging two system message logging databases, follow these guidelines:

- Be aware that the merged database is a union of the existing and received database for each switch in the fabric.
- Verify that the merged database will only have a maximum of three system message logging servers.



Caution

If the merged database contains more than three servers, the merge will fail.

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying System Message Logging Information

Use the **show logging** command to display the current system message logging configuration. See Examples 36-1 to 36-10.

Example 36-1 *Displays Current System Message Logging*

```
switch# show logging
Logging console:          enabled (Severity: critical)
Logging monitor:          enabled (Severity: debugging)
Logging linecard:         enabled (Severity: debugging)
Logging server:           enabled
{172.20.102.34}
    server severity:      debugging
    server facility:      local7
{10.77.202.88}
    server severity:      debugging
    server facility:      local7
{10.77.202.149}
    server severity:      debugging
    server facility:      local7
Logging logfile:          enabled
Name - messages: Severity - debugging Size - 4194304
Facility      Default Severity      Current Session Severity
-----
kern          6
user          3
mail          3
daemon       7
auth          0
syslog        3
lpr           3
news          3
uucp          3
cron          3
authpriv      3
ftp           3
local0        3
local1        3
local2        3
local3        3
local4        3
local5        3
local6        3
local7        3
vsan          2
fspf          3
fcdomain      2
module        5
sysmgr        3
zone          2
vni           2
ipconf        2
ipfc          2
xbar          3
fcns          2
fcs           2
acl           2
tlport        2
port          5
flogi         2
port_channel  5
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```

wwn                3                3
fcc                2                2
qos                3                3
vrrp_cfg           2                2
ntp                2                2
platform           5                5
vrrp_eng           2                2
callhome           2                2
mcast              2                2
rdl                2                2
rscn               2                2
bootvar            5                2
securityd          2                2
vhbad              2                2
rib                2                2
vshd               5                5
0(emergencies)     1(alerts)        2(critical)
3(errors)          4(warnings)      5(notifications)
6(information)     7(debugging)

```

```

Feb 14 09:50:57 excal-113 %TTYD-6-TTYD_MISC: TTYD TTYD started
Feb 14 09:50:58 excal-113 %DAEMON-6-SYSTEM_MSG: precision = 8 usec
...

```

Use the **show logging nvram** command to view the log messages saved in NVRAM. Only log messages with a severity level of critical and below (levels 0, 1, and 2) are saved in NVRAM.

Example 36-2 Displays NVRM Log Contents

```

switch# show logging nvram
Jul 16 20:36:46 172.22.91.204 %KERN-2-SYSTEM_MSG: unable to alloc and fill in a
new mtsbuf (pid=2209, ret_val = -105)
Jul 16 20:36:46 172.22.91.204 %KERN-2-SYSTEM_MSG: unable to alloc and fill in a
new mtsbuf (pid=2199, ret_val = -105)
Jul 16 20:36:46 172.22.91.204 %KERN-2-SYSTEM_MSG: unable to alloc and fill in a
new mtsbuf (pid=2213, ret_val = -105)
Jul 16 20:36:46 172.22.91.204 %KERN-2-SYSTEM_MSG: unable to alloc and fill in a
new mtsbuf (pid=2213, ret_val = -105)
...

```

Example 36-3 Displays the Log File

```

switch# show logging logfile
Jul 16 21:06:50 %DAEMON-3-SYSTEM_MSG: Un-parsable frequency in /mnt/pss/ntp.drift
Jul 16 21:06:56 %DAEMON-3-SYSTEM_MSG: snmpd:snmp_open_debug_cfg: no snmp_saved_dbg_uri ;
Jul 16 21:06:58 172.22.91.204 %PORT-5-IF_UP: Interface mgmt0 is up
Jul 16 21:06:58 172.22.91.204 %MODULE-5-ACTIVE_SUP_OK: Supervisor 5 is active
...

```

Example 36-4 Displays Console Logging Status

```

switch# show logging console
Logging console:                enabled (Severity: notifications)

```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 36-5 Displays Logging Facility

```
switch# show logging level
Facility           Default Severity    Current Session Severity
-----
kern                6                    6
user                3                    3
mail                3                    3
daemon              7                    7
auth                0                    7
syslog              3                    3
lpr                 3                    3
news                3                    3
uucp                3                    3
cron                3                    3
authpriv            3                    7
ftp                 3                    3
local0              3                    3
local1              3                    3
local2              3                    3
local3              3                    3
local4              3                    3
local5              3                    3
local6              3                    3
local7              3                    3
vsan                2                    2
fspf                3                    3
fcdomain            2                    2
module              5                    5
sysmgr              3                    3
zone                2                    2
vni                 2                    2
ipconf              2                    2
ipfc                2                    2
xbar                3                    3
fcns                2                    2
fcs                 2                    2
acl                 2                    2
tlport              2                    2
port                5                    5
flogi               2                    2
port_channel        5                    5
wnn                 3                    3
fcc                 2                    2
qos                 3                    3
vrrp_cfg            2                    2
ntp                 2                    2
platform            5                    5
vrrp_eng            2                    2
callhome            2                    2
mcast               2                    2
rdl                 2                    2
rscn                2                    2
bootvar             5                    2
securityd           2                    2
vhbad               2                    2
rib                 2                    2
vshd                5                    5
0(emergencies)      1(alerts)            2(critical)
3(errors)            4(warnings)          5(notifications)
6(information)      7(debugging)
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 36-6 Displays Logging Information

```
switch# show logging info
Logging console:          enabled (Severity: critical)
Logging monitor:          enabled (Severity: debugging)
Logging linecard:         enabled (Severity: debugging)
Logging server:           enabled
{172.20.102.34}
    server severity:      debugging
    server facility:      local7
{10.77.202.88}
    server severity:      debugging
    server facility:      local7
{10.77.202.149}
    server severity:      debugging
    server facility:      local7
Logging logfile:          enabled
    Name - messages: Severity - debugging Size - 4194304
```

| Facility | Default Severity | Current Session Severity |
|--------------|------------------|--------------------------|
| ----- | ----- | ----- |
| kern | 6 | 6 |
| user | 3 | 3 |
| mail | 3 | 3 |
| daemon | 7 | 7 |
| auth | 0 | 7 |
| syslog | 3 | 3 |
| lpr | 3 | 3 |
| news | 3 | 3 |
| uucp | 3 | 3 |
| cron | 3 | 3 |
| authpriv | 3 | 7 |
| ftp | 3 | 3 |
| local0 | 3 | 3 |
| local1 | 3 | 3 |
| local2 | 3 | 3 |
| local3 | 3 | 3 |
| local4 | 3 | 3 |
| local5 | 3 | 3 |
| local6 | 3 | 3 |
| local7 | 3 | 3 |
| vsan | 2 | 2 |
| fspf | 3 | 3 |
| fcdomain | 2 | 2 |
| module | 5 | 5 |
| sysmgr | 3 | 3 |
| zone | 2 | 2 |
| vni | 2 | 2 |
| ipconf | 2 | 2 |
| ipfc | 2 | 2 |
| xbar | 3 | 3 |
| fcns | 2 | 2 |
| fcs | 2 | 2 |
| acl | 2 | 2 |
| tlport | 2 | 2 |
| port | 5 | 5 |
| flogi | 2 | 2 |
| port_channel | 5 | 5 |
| wwn | 3 | 3 |
| fcc | 2 | 2 |
| qos | 3 | 3 |
| vrrp_cfg | 2 | 2 |
| ntp | 2 | 2 |
| platform | 5 | 5 |
| vrrp_eng | 2 | 2 |

Send documentation comments to mdsfeedback-doc@cisco.com.

```

callhome                2                2
mcast                   2                2
rdl                     2                2
rscn                    2                2
bootvar                 5                2
securityd               2                2
vhbad                   2                2
rib                     2                2
vshd                    5                5
0(emergencies)          1(alerts)        2(critical)
3(errors)                4(warnings)      5(notifications)
6(information)          7(debugging)

```

Example 36-7 Displays Last Few Lines of a Log File

```

switch# show logging last 2
Nov 8 16:48:04 excal-113 %LOG_VSHD-5-VSHD_SYSLOG_CONFIG_I: Configuring console from pts/1
(171.71.58.56)
Nov 8 17:44:09 excal-113 %LOG_VSHD-5-VSHD_SYSLOG_CONFIG_I: Configuring console from pts/0
(171.71.58.72)

```

Example 36-8 Displays Switching Module Logging Status

```

switch# show logging module
Logging linecard:                enabled (Severity: debugging)

```

Example 36-9 Displays Monitor Logging Status

```

switch# show logging monitor
Logging monitor:                enabled (Severity: information)

```

Example 36-10 Displays Server Information

```

switch# show logging server
Logging server:                enabled
{172.22.95.167}
    server severity:           debugging
    server facility:           local7
{172.22.92.58}
    server severity:           debugging
    server facility:           local7

```

Default Settings

Table 36-4 lists the default settings for system message logging.

Table 36-4 Default System Message Log Setting

| Parameters | Default |
|---|---|
| System message logging to the console | Enabled for messages at the critical severity level. |
| System message logging to Telnet sessions | Disabled. |
| Logging file size | 4194304. |
| Log file name | Message (change to a name with up to 200 characters). |
| Logging server | Disabled. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 36-4 **Default System Message Log Setting (continued)**

| Parameters | Default |
|--------------------------|-----------------|
| Syslog server IP address | Not configured. |
| Number of servers | Three servers. |
| Server facility | Local 7. |



Discovering SCSI Targets

This chapter describes the SCSI LUN discovery feature provided in switches in the Cisco MDS 9000 Family. It includes the following sections:

- [About SCSI LUN Discovery, page 37-1](#)
- [Starting SCSI LUN Discovery, page 37-2](#)
- [Initiating Customized Discovery, page 37-2](#)
- [Displaying SCSI LUN Information, page 37-3](#)

About SCSI LUN Discovery

Small Computer System Interface (SCSI) targets include disks, tapes, and other storage devices. These targets do not register logical unit numbers (LUNs) with the name server.

The name server requires LUN information for the following reasons:

- To display LUN storage device information so an NMS can access this information.
- To report device capacity, serial number, and device ID information.
- To register the initiator and target features with the name server.

The SCSI LUN discovery feature uses the local domain controller Fibre Channel address. It uses the local domain controller as the source FC ID, and performs SCSI INQUIRY, REPORT LUNS, and READ CAPACITY commands on SCSI devices.

The SCSI LUN discovery feature is initiated on demand, through CLI or SNMP. This information is also synchronized with neighboring switches, if those switches belong to the Cisco MDS 9000 Family.

Send documentation comments to mdsfeedback-doc@cisco.com.

Starting SCSI LUN Discovery

SCSI LUN discovery is done on demand.

Only Nx ports present in the name server database and that are registered as FC4 Type = SCSI_FCP are discovered.

To begin SCSI LUN discovery, follow this step:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# discover scsi-target local os all discovery started | Discovers local SCSI targets for all operating systems (OS). The operating system options are aix , all , hpux , linux , solaris , or windows |
| | switch# discover scsi-target remote os aix discovery started | Discovers remote SCSI targets assigned to the AIX OS. |
| | switch# discover scsi-target vsan 1 fcid 0x9c03d6 discover scsi-target vsan 1 fcid 0x9c03d6 VSAN: 1 FCID: 0x9c03d6 PWWN: 00:00:00:00:00:00:00:00 PRLI RSP: 0x01 SPARM: 0x0012 SCSI TYPE: 0 NLUNS: 1 Vendor: Company 4 Model: ST318203FC Rev: 0004 Other: 00:00:02:32:8b:00:50:0a | Discovers SCSI targets for the specified VSAN (1) and FC ID (0x9c03d6). |
| | switch# discover scsi-target custom-list os linux discovery started | Discovers SCSI targets from the customized list assigned to the Linux OS. |

Initiating Customized Discovery

Customized discovery consists of a list of VSAN and domain pairs that are selectively configured to initiate a discovery. Use the **custom-list** option to initiate this discovery. The domain ID is a number from 0 to 255 in decimal or a number from 0x0 to 0xFF in hex.

To initiate a customized discovery, follow this step:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# discover custom-list add vsan 1 domain 0X123456 | Adds the specified entry to the custom list. |
| | switch# discover custom-list delete vsan 1 domain 0X123456 | Deletes the specified domain ID from the custom list. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying SCSI LUN Information

Use the **show scsi-target** and **show fcns database** commands to display the results of the discovery. See Examples 37-1 to 37-8.

Example 37-1 Displays the Discovered Targets

```
switch# show scsi-target status
discovery completed
```



Note

This command takes several minutes to complete, especially if the fabric is large or if several devices are slow to respond.

Example 37-2 Displays the FCNS Database

```
switch# show fcns database
172.22.91.115# show fcns database
```

VSAN 1:

| FCID | TYPE | PWWN | (VENDOR) | FC4-TYPE:FEATURE |
|----------|------|-------------------------|----------|------------------|
| 0xeb0000 | N | 21:01:00:e0:8b:2a:f6:54 | (Qlogic) | scsi-fcp:init |
| 0xeb0201 | NL | 10:00:00:00:c9:32:8d:76 | (Emulex) | scsi-fcp:init |

Total number of entries = 2

VSAN 7:

| FCID | TYPE | PWWN | (VENDOR) | FC4-TYPE:FEATURE |
|----------|------|-------------------------|-----------|------------------|
| 0xed0001 | NL | 21:00:00:04:cf:fb:42:f8 | (Seagate) | scsi-fcp:target |

Total number of entries = 1

VSAN 2002:

| FCID | TYPE | PWWN | (VENDOR) | FC4-TYPE:FEATURE |
|----------|------|-------------------------|----------|------------------|
| 0xcafe00 | N | 20:03:00:05:30:00:2a:20 | (Cisco) | FICON:CUP |

Total number of entries = 1

Example 37-3 Displays the Discovered Target Disks

```
switch# show scsi-target disk
```

| VSAN | FCID | PWWN | VENDOR | MODEL | REV |
|------|----------|-------------------------|-----------|-----------------|------|
| 1 | 0x9c03d6 | 21:00:00:20:37:46:78:97 | Company 4 | ST318203FC | 0004 |
| 1 | 0x9c03d9 | 21:00:00:20:37:5b:cf:b9 | Company 4 | ST318203FC | 0004 |
| 1 | 0x9c03da | 21:00:00:20:37:18:6f:90 | Company 4 | ST318203FC | 0004 |
| 1 | 0x9c03dc | 21:00:00:20:37:5a:5b:27 | Company 4 | ST318203FC | 0004 |
| 1 | 0x9c03e0 | 21:00:00:20:37:36:0b:4d | Company 4 | ST318203FC | 0004 |
| 1 | 0x9c03e1 | 21:00:00:20:37:39:90:6a | Company 4 | ST318203 CLAR18 | 3844 |
| 1 | 0x9c03e2 | 21:00:00:20:37:18:d2:45 | Company 4 | ST318203 CLAR18 | 3844 |
| 1 | 0x9c03e4 | 21:00:00:20:37:6b:d7:18 | Company 4 | ST318203 CLAR18 | 3844 |
| 1 | 0x9c03e8 | 21:00:00:20:37:38:a7:c1 | Company 4 | ST318203FC | 0004 |
| 1 | 0x9c03ef | 21:00:00:20:37:18:17:d2 | Company 4 | ST318203FC | 0004 |

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 37-4 Displays the Discovered LUNs for All OSs

```
switch# show scsi-target lun os all
ST336607FC from SEAGATE (Rev 0006)
FCID is 0xed0001 in VSAN 7, PWWN is 21:00:00:04:cf:fb:42:f8
-----
OS   LUN   Capacity Status   Serial Number   Device-Id
      (MB)
-----
WIN 0x0   36704   Online  3JA1B9QA00007338 C:1 A:0 T:3 20:00:00:04:cf:fb:42:f8
AIX 0x0   36704   Online  3JA1B9QA00007338 C:1 A:0 T:3 20:00:00:04:cf:fb:42:f8
SOL 0x0   36704   Online  3JA1B9QA00007338 C:1 A:0 T:3 20:00:00:04:cf:fb:42:f8
LIN 0x0   36704   Online  3JA1B9QA00007338 C:1 A:0 T:3 20:00:00:04:cf:fb:42:f8
HP  0x0   36704   Online  3JA1B9QA00007338 C:1 A:0 T:3 20:00:00:04:cf:fb:42:f8
```

Example 37-5 Displays the Discovered LUNs for the Solaris OS

```
switch# show scsi-target lun os solaris
ST336607FC from SEAGATE (Rev 0006)
FCID is 0xed0001 in VSAN 7, PWWN is 21:00:00:04:cf:fb:42:f8
-----
OS   LUN   Capacity Status   Serial Number   Device-Id
      (MB)
-----
SOL 0x0   36704   Online  3JA1B9QA00007338 C:1 A:0 T:3 20:00:00:04:cf:fb:42:f8
```

The following command displays the port WWN that is assigned to each OS (Windows, AIX, Solaris, Linux, or HP/UX)

Example 37-6 Displays the pWWNs for each OS

```
switch# show scsi-target pwwn
-----
OS      PWWN
-----
WIN     24:91:00:05:30:00:2a:1e
AIX     24:92:00:05:30:00:2a:1e
SOL     24:93:00:05:30:00:2a:1e
LIN     24:94:00:05:30:00:2a:1e
HP      24:95:00:05:30:00:2a:1e
```

Example 37-7 Displays Customized Discovered Targets

```
switch# show scsi-target custom-list
-----
VSAN    DOMAIN
-----
1        56
```

Use the **show scsi-target auto-poll** command to verify automatic discovery of SCSI targets that come online. The internal uuid number indicates that a CSM or an IPS module is in the chassis.

Example 37-8 Displays Customized Discovered Targets

```
switch# show scsi-target auto-poll
auto-polling is enabled, poll_start:0 poll_count:1 poll_type:0
USERS OF AUTO POLLING
-----
uuid:54
```



Monitoring Network Traffic Using SPAN

This chapter describes the Switched Port Analyzer (SPAN) features provided in switches in the Cisco MDS 9000 Family. It includes the following sections:

- [About SPAN, page 38-2](#)
- [SPAN Sources, page 38-3](#)
- [SPAN Sessions, page 38-5](#)
- [Specifying Filters, page 38-6](#)
- [SD Port Characteristics, page 38-6](#)
- [Configuring SPAN, page 38-7](#)
- [Monitoring Traffic Using Fibre Channel Analyzers, page 38-10](#)
- [Displaying SPAN Information, page 38-13](#)
- [Remote SPAN, page 38-15](#)
- [Default SPAN and RSPAN Settings, page 38-29](#)

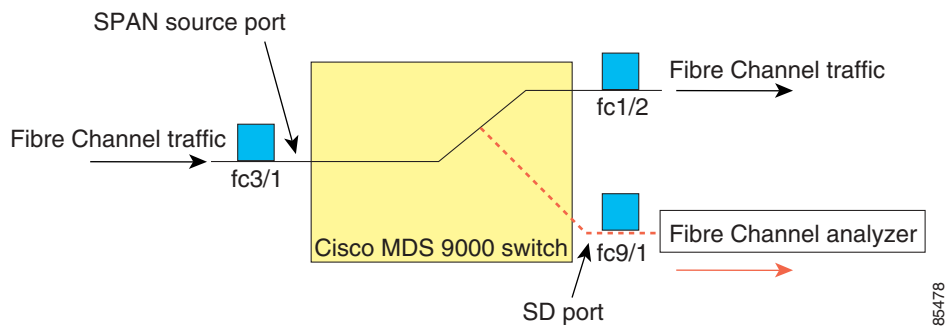
Send documentation comments to mdsfeedback-doc@cisco.com.

About SPAN

The SPAN feature is specific to switches in the Cisco MDS 9000 Family. It monitors network traffic through a Fibre Channel interface. Traffic through any Fibre Channel interface can be replicated to a special port called the SPAN destination port (SD port). Any Fibre Channel port in a switch can be configured as an SD port. Once an interface is in SD port mode, it cannot be used for normal data traffic. You can attach a Fibre Channel Analyzer to the SD port to monitor SPAN traffic (see [“Configuring a Fabric Analyzer”](#) section on page 39-8).

SD ports do not receive frames, they merely transmit a copy of the SPAN source traffic. The SPAN feature is non-intrusive and does not affect switching of network traffic for any SPAN source ports (see [Figure 38-1](#)).

Figure 38-1 *SPAN Transmission*




Send documentation comments to mdsfeedback-doc@cisco.com.

SPAN Sources

SPAN sources refer to the interfaces from which traffic can be monitored. You can also specify VSAN as a SPAN source, in which case, all supported interfaces in the specified VSAN are included as SPAN sources. You can choose the SPAN traffic in the ingress direction, the egress direction, or both directions for any source interface:


- Ingress source (Rx)—Traffic entering the switch fabric through this source interface is *spanned* or copied to the SD port (see Figure 38-2).

Figure 38-2 SPAN Traffic from the Ingress Direction



- Egress source (Tx)—Traffic exiting the switch fabric through this source interface is spanned or copied to the SD port (see Figure 38-3).

Figure 38-3 SPAN Traffic from Egress Direction



IPS Source Ports

As of Cisco MDS SAN-OS Release 1.3 SPAN capabilities are also available on the IP Storage Services (IPS) module. The SPAN feature is only implemented on the FCIP and iSCSI virtual Fibre Channel port interfaces, not the physical Gigabit Ethernet ports. You can configure SPAN for ingress traffic, egress traffic, or traffic in both directions for all eight iSCSI and 24 FCIP interfaces that are available in the IPS module.



Note

You can configure SPAN for Ethernet traffic using Cisco switches or routers connected to the Cisco MDS 9000 Family IPS modules.

Send documentation comments to mdsfeedback-doc@cisco.com.

CSM Source Ports

As of Cisco MDS SAN-OS Release 1.3 SPAN capabilities are also available on the Caching Services Module (CSM).

Refer to the *Cisco MDS 9000 Family SAN Volume Controller Configuration Guide* for more information.

Allowed Source Interface Types

The SPAN feature is available for the following interface types:

- Physical ports such as F ports, FL ports, TE ports, E ports, and TL ports.
- Interface sup-fc0 (traffic to and from the supervisor):
 - The Fibre Channel traffic from the supervisor module to the switch fabric through the sup-fc0 interface is called ingress traffic. It is spanned when sup-fc0 is chosen as an ingress source port.
 - The Fibre Channel traffic from the switch fabric to the supervisor module through the sup-fc0 interface is called egress traffic. It is spanned when sup-fc0 is chosen as an egress source port.
- PortChannels
 - All ports in the PortChannel are included and spanned as sources.
 - You cannot specify individual ports in a PortChannel as SPAN sources. Previously configured SPAN-specific interface information is discarded.
- IPS module specific Fibre Channel interfaces.
 - iSCSI interfaces
 - FCIP interfaces

VSAN as a Source

When a VSAN as a source is specified, then all physical ports and PortChannels in that VSAN are included as SPAN sources. A TE port is included only when the port VSAN of the TE port matches the source VSAN. A TE port is excluded even if the configured allowed VSAN list may have the source VSAN, but the port VSAN is different.

You cannot configure source interfaces (physical interfaces, PortChannels, or sup-fc interfaces) and source VSANs in the same SPAN session.

Guidelines to Configure VSANs as a Source

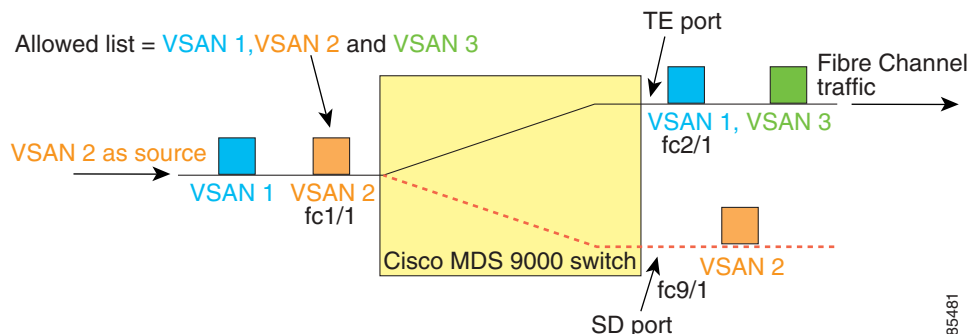
The following guidelines apply when configuring VSANs as a source:

- Traffic on all interfaces included in a source VSAN is spanned only in the ingress direction.
- If a VSAN is specified as a source, you cannot perform interface-level SPAN configuration on the interfaces that are included in the VSAN. Previously configured SPAN-specific interface information is discarded.
- If an interface in a VSAN is configured as a source, you cannot configure that VSAN as a source. You must first remove the existing SPAN configurations on such interfaces before configuring VSAN as a source.

Send documentation comments to mdsfeedback-doc@cisco.com.

- Interfaces are only included as sources when the port VSAN matches the source VSAN. [Figure 38-4](#) displays a configuration using VSAN 2 as a source:
 - All ports in the switch are in VSAN 1 except fc1/1.
 - Interface fc1/1 is the TE port with port VSAN 2. VSANs 1, 2, and 3 are configured in the allowed list.
 - VSAN 1 and VSAN 2 are configured as SPAN sources.

Figure 38-4 VSAN As a Source



For this configuration, the following apply:

- VSAN 2 as a source includes only the TE port fc1/1 that has port VSAN 2.
- VSAN 1 as a source does not include the TE port fc1/1 as the port VSAN does not match VSAN 1.

See the “Configuring an Allowed-Active List of VSANs” section on page 13-5 or the “VSAN Membership” section on page 10-6.

SPAN Sessions

Each SPAN session represents an association of one destination with a set of source(s) along with various other parameters that you specify to monitor the network traffic. One destination can be used by one or more SPAN sessions. You can configure up to 16 SPAN sessions in a switch. Each session can have several source ports and one destination port.

To activate any SPAN session, at least one source and the SD port must be up and functioning. Otherwise, traffic is not directed to the SD port.



Tip

A source can be shared by two sessions, however, each session must be in a different direction—one ingress and one egress.

To temporarily deactivate (suspend) any SPAN session, use the **suspend** command in the SPAN submode. The traffic monitoring is stopped during this time. You can reactivate the SPAN session using the **no suspend** command.

Send documentation comments to mdsfeedback-doc@cisco.com.

Specifying Filters

You can perform VSAN-based filtering to selectively monitor network traffic on specified VSANs. You can apply this VSAN filter to all sources in a session (see [Figure 38-4](#)). Only VSANs present in the filter are spanned.

You can specify session VSAN filters that are applied to all sources in the specified session. These filters are bidirectional and apply to all sources configured in the session.

Guidelines to Specifying Filters

The following guidelines apply to SPAN filters:

- PortChannel configurations are applied to all ports in the PortChannel.
- If no filters are specified, the traffic from all active VSANs for that interface is spanned by default.
- While you can specify arbitrary VSAN filters in a session, traffic can only be monitored on the port VSAN or on allowed-active VSANs in that interface.

SD Port Characteristics

An SD port has the following characteristics:

- Ignores BB_credits.
- Allows data traffic only in the egress (Tx) direction.
- Does not require a device or an analyzer to be physically connected.
- Supports only 1 Gbps or 2 Gbps speeds. The auto speed option is not allowed.
- Multiple sessions can share the same destination ports.
- If the SD port is shut down, all shared sessions stop generating SPAN traffic.
- The outgoing frames can be encapsulated in Extended Inter-Switch Link (EISL) format.
- The SD port does not have a port VSAN.
- SD ports cannot be configured using Advanced Services Modules (ASMs) or Storage Services Modules (SSMs).
- The port mode cannot be changed if it is being used for a SPAN session.



Note

If you need to change an SD port mode to another port mode, first remove the SD port from all sessions and then change the port mode using the **switchport mode** command.

Guidelines to Configure SPAN

The following guidelines apply for SPAN configurations:

- You can configure up to 16 SPAN sessions with multiple ingress (Rx) sources.
- You can configure a maximum of three SPAN sessions with one egress (Tx) port.

Send documentation comments to mdsfeedback-doc@cisco.com.

- In a 32-port switching module, you must configure the same session in all four ports in one port group (unit). If you wish, you can also configure only two or three ports in this unit (see the “[32-Port Configuration Guidelines](#)” section on page 12-8).
- SPAN frames are dropped if the sum of the bandwidth of the sources exceeds the speed of the destination port.
- Frames dropped by a source port are not spanned.

Configuring SPAN

To monitor network traffic using SD ports, follow these steps:

- | | |
|---------------|---|
| Step 1 | Configure the SD port. |
| Step 2 | Attach the SD port to a specific SPAN session. |
| Step 3 | Monitor network traffic by adding source interfaces to the session. |

To configure an SD port for SPAN monitoring, follow these steps:

| | Command | Purpose |
|---------------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc9/1 | Configures the specified interface. |
| Step 3 | switch(config-if)# switchport mode SD | Configures the SD port mode for interface fc2/1. |
| Step 4 | switch(config-if)# switchport speed 1000 | Configures the SD port speed to 1000 Mbps. |
| Step 5 | switch(config-if)# no shutdown | Enables traffic flow through this interface. |

To configure a SPAN session, follow these steps:

| | Command | Purpose |
|---------------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# span session 1 switch(config-span)# | Configures the specified SPAN session (1). If the session does not exist, it is created. |
| | switch(config)# no span session 1 | Deletes the specified SPAN session (1). |
| Step 3 | switch(config-span)# destination interface fc9/1 | Configures the specified destination interface (fc 9/1) in a session. |
| | switch(config-span)# no destination interface fc9/1 | Removes the specified destination interface (fc 9/1). |
| Step 4 | switch(config-span)# source interface fc7/1 | Configures the source (fc7/1) interface in both directions. |
| | switch(config-span)# no source interface fc7/1 | Removes the specified destination interface (fc 7/1) from this session. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|---|--|
| Step 5 | switch(config-span)# source interface sup-fc0 | Configures the source interface (sup-fc0) in the session. |
| | switch(config-span)# source interface fc1/5 - 6, fc2/1 -3 | Configures the specified interface ranges in the session. |
| | switch(config-span)# source vsan 1-2 | Configures source VSANs 1 and 2 in the session. |
| | switch(config-span)# source interface port-channel 1 | Configures the source PortChannel (port-channel 1). |
| | switch(config-span)# source interface fcip 51 | Configures the source FCIP interface in the session. |
| | switch(config-span)# source interface iscsi 4/1 | Configures the source iSCSI interface in the session. |
| | switch(config-span)# source interface svc1/1 tx traffic-type initiator | Configures the source SVC interface in the Tx direction for an initiator traffic type. |
| | switch(config-span)# no source interface port-channel 1 | Deletes the specified source interface (port-channel 1). |
| Step 6 | switch(config-span)# suspend | Suspends the session. |
| | switch(config-span)# no suspend | Reactivates the session. |

To configure a SPAN filter, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# span session 1 switch(config-span)# | Configures the specified session (1). |
| Step 3 | switch(config-span)# source interface fc9/1 tx | Configures the source fc9/1 interface in the egress (Tx) direction. |
| | switch(config-span)# source filter vsan 1-2 | Configures VSANs 1 and 2 as session filters. |
| | switch(config-span)# source interface fc7/1 rx | Configures the source fc7/1 interface in the ingress (Rx) direction. |

Encapsulating Frames

The frame encapsulation feature is disabled by default. If you enable the encapsulation feature, all outgoing frames are encapsulated.

The **switchport encap eisl** command only applies to SD port interfaces. If encapsulation is enabled, you see a new line (Encapsulation is eisl) in the **show interface SD_port_interface** command output.

To encapsulate outgoing frames (optional), follow these steps:

| | Command | Purpose |
|--------|---|-------------------------------------|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc9/32 | Configures the specified interface. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|--|--|
| Step 3 | switch(config-if)# switchport mode SD | Configures the SD port-mode for interface fc2/1. |
| Step 4 | switch(config-if)# switchport encap eisl | Enables the encapsulation option for this SD port. |
| | switch(config-if)# no switchport encap eisl | Disables the encapsulation option and reverts the switch to factory default. |

SPAN Conversion Behavior

As of Cisco MDS SAN-OS Release 1.1(1), SPAN features (configured in any prior release) are converted as follows:

- If source interfaces and source VSANs are configured in a given session, then all the source VSANs are removed from that session.

For example,

Before Cisco MDS SAN-OS Release 1.0(4):

```
Session 1 (active)
  Destination is fc1/9
  No session filters configured
  Ingress (rx) sources are
    vsans 10-11
    fc1/3,
  Egress (tx) sources are
    fc1/3,
```

Once upgraded to Cisco MDS SAN-OS Release 1.1(1):

```
Session 1 (active)
  Destination is fc1/9
  No session filters configured
  Ingress (rx) sources are
    fc1/3,
  Egress (tx) sources are
    fc1/3,
```

Session 1 had both source interfaces and source VSANs before the upgrade. After the upgrade, the source VSANs were removed (rule 1).

- If interface level VSAN filters are configured in source interfaces, then the source interfaces are also removed from the session. If this interface is configured in both directions, it is removed from both directions.

For example,

Before Cisco MDS SAN-OS Release 1.0(4):

```
Session 2 (active)
  Destination is fc1/9
  No session filters configured
  Ingress (rx) sources are
    vsans 12
    fc1/6 (vsan 1-20),
  Egress (tx) sources are
    fc1/6 (vsan 1-20),
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Once upgraded to Cisco MDS SAN-OS Release 1.1(1):

```
Session 2 (inactive as no active sources)
Destination is fc1/9
No session filters configured
No ingress (rx) sources
No egress (tx) sources
```



Note

The deprecated configurations are removed from persistent memory once a switchover or a new startup configuration is implemented.

Session 2 had a source VSAN 12 and a source interface fc1/6 with VSAN filters specified in Cisco MDS SAN-OS Release 1.0(4). When upgraded to Cisco MDS SAN-OS Release 1.1(1) the following changes are made:

- The source VSAN (VSAN 12) is removed (rule 1).
- The source interface fc1/6 had VSAN filters specified—it is also removed (rule 2).

Monitoring Traffic Using Fibre Channel Analyzers


You can use SPAN to monitor traffic on an interface without any traffic disruption. This feature is specially useful in troubleshooting scenarios where traffic disruption changes the problem environment and makes it difficult to reproduce the problem.

Without SPAN

You can monitor traffic using interface fc1/1 in a Cisco MDS 9000 Family switch that is connected to another switch or host. You need to physically connect a Fibre Channel analyzer between the switch and the storage device to analyze the traffic through interface fc1/1 as shown in [Figure 38-5](#).

Figure 38-5 Fibre Channel Analyzer Usage Without SPAN

FC Analyzer usage without SPAN



Send documentation comments to mdsfeedback-doc@cisco.com.

This type of connection has the following limitations:

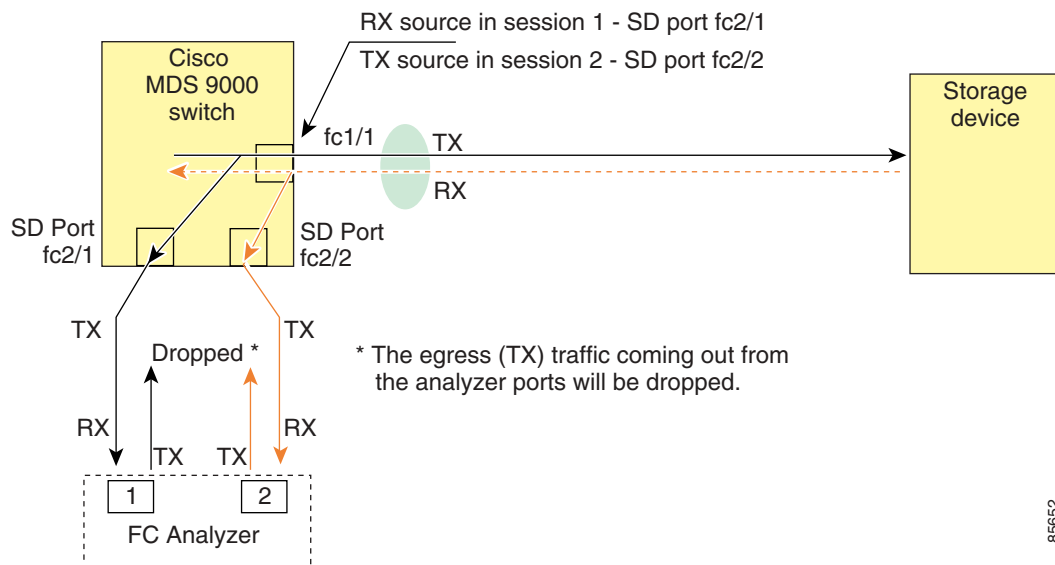
- It requires you to physically insert the FC analyzer between the two network devices.
- It disrupts traffic when the Fibre Channel analyzer is physically connected.
- The analyzer captures data only on the Rx links in both port 1 and port 2. Port 1 captures traffic exiting interface fc1/1 and port 2 captures ingress traffic into interface fc1/1.

With SPAN

Using SPAN you can capture the same traffic scenario shown in [Figure 38-5](#) without any traffic disruption. The Fibre Channel analyzer uses the ingress (Rx) link at port 1 to capture all the frames going out of the interface fc1/1. It uses the ingress link at port 2 to capture all the ingress traffic on interface fc1/1.

Using SPAN you can monitor ingress traffic on fc1/1 at SD port fc2/2 and egress traffic on SD port fc2/1. This traffic is seamlessly captured by the FC analyzer as shown in [Figure 38-6](#).

Figure 38-6 Fibre Channel Analyzer Using SPAN



85652

Configuring Analyzers Using SPAN

To configure Fibre Channel Analyzers using SPAN for the example in [Figure 38-6](#), follow these steps:

- Step 1** Configure SPAN on interface fc1/1 in the ingress (Rx) direction to send traffic on SD port fc2/1 using session 1.
- Step 2** Configure SPAN on interface fc1/1 in the egress (Tx) direction to send traffic on SD port fc2/2 using session 2.
- Step 3** Physically connect fc2/1 to port 1 on the Fibre Channel analyzer.
- Step 4** Physically connect fc2/2 to port 2 on the Fibre Channel analyzer.

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure SPAN on the source and destination interfaces, follow these steps:

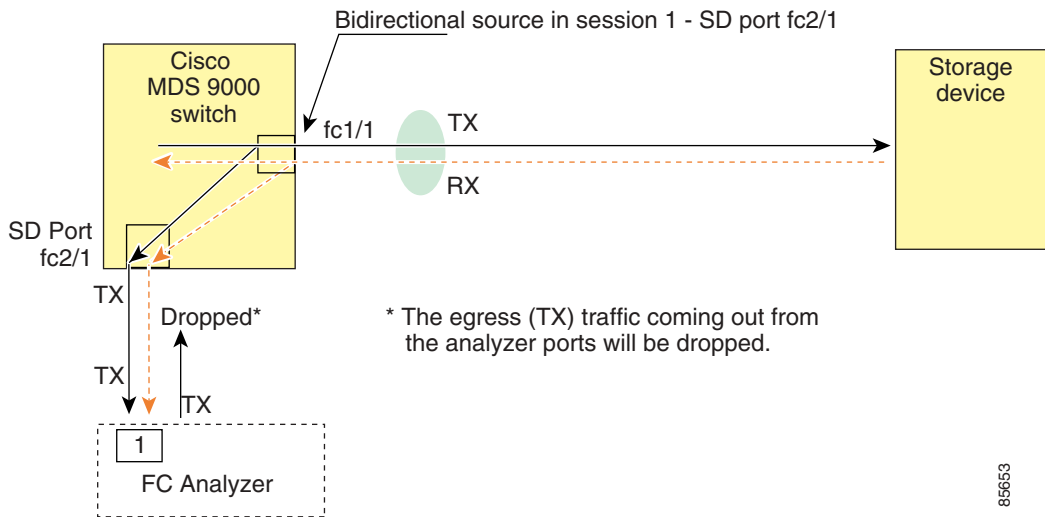
| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# confi g t | Enters configuration mode. |
| Step 2 | switch(config)# span session 1 switch(config-span)# | Creates the SPAN session 1. |
| Step 3 | switch(config-span)## destination interface fc2/1 | Configures the destination interface fc2/1. |
| Step 4 | switch(config-span)# source interface fc1/1 rx | Configures the source interface fc1/1 in the ingress direction. |
| Step 5 | switch(config)# span session 2 switch(config-span)# | Creates the SPAN session 2. |
| Step 6 | switch(config-span)## destination interface fc2/2 | Configures the destination interface fc2/2. |
| Step 7 | switch(config-span)# source interface fc1/1 tx | Configures the source interface fc1/1 in the egress direction. |

Single SD Port to Monitor Traffic

You do not need to use two SD ports to monitor bidirectional traffic on any interface as shown in Figure 38-6. You can use one SD port and one FC analyzer port by monitoring traffic on the interface at the same SD port fc2/1.

Figure 38-7 shows a SPAN setup where one session with destination port fc2/1 and source interface fc1/1 is used to capture traffic in both ingress and egress direction. This setup is more advantageous and cost effective than the setup shown in Figure 38-6—it uses one SD port and one port on the analyzer, instead of using a full, two-port analyzer.

Figure 38-7 Fibre Channel Analyzer Using a Single SD Port



To use this setup, the analyzer should have the capability of distinguishing ingress and egress traffic for all captured frames.

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure SPAN on a single SD port, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# span session 1 switch(config-span)# | Creates the SPAN session 1. |
| Step 3 | switch(config-span)## destination interface fc2/1 | Configures the destination interface fc2/1. |
| Step 4 | switch(config-span)# source interface fc1/1 | Configures the source interface fc1/1 on the same SD port. |

Displaying SPAN Information

Use the **show span** command to display configured SPAN information. See Examples 38-1 to 38-4.

Example 38-1 Displays SPAN Sessions in a Brief Format

```
switch# show span session brief
-----
Session  Admin      Oper      Destination
         State        State      Interface
-----
 7       no suspend    active    fc2/7
 1       suspend     inactive  not configured
 2       no suspend    inactive  fc3/1
```

Example 38-2 Displays a Specific SPAN Session in Detail

```
switch# show span session 7
Session 7 (active)
  Destination is fc2/7
  No session filters configured
  No ingress (rx) sources
  Egress (tx) sources are
    port-channel 7,
```

Example 38-3 Displays ALL SPAN Sessions

```
switch# show span session
Session 1 (inactive as no destination)
Destination is not specified
  Session filter vsans are 1
  No ingress (rx) sources
  No egress (tx) sources
Session 2 (active)
  Destination is fc9/5
  No session filters configured
  Ingress (rx) sources are
    vsans 1
  No egress (tx) sources
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
Session 3 (admin suspended)
  Destination is not configured
  Session filter vsans are 1-20
  Ingress (rx) sources are
    fc3/2, fc3/3, fc3/4, fcip 51,
    port-channel 2, sup-fc0,
  Egress (tx) sources are
    fc3/2, fc3/3, fc3/4, sup-fc0,
```

Example 38-4 Displays an SD Port Interface with Encapsulation Enabled

```
switch# show int fc9/32
fc9/32 is up
  Hardware is Fibre Channel
  Port WWN is 22:20:00:05:30:00:49:5e
  Admin port mode is SD
  Port mode is SD
  Port vsan is 1
  Speed is 1 Gbps
  Receive Buffer Size is 2112
  Encapsulation is eisl <----- Displays the enabled encapsulation status
  Beacon is turned off
  5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    0 frames input, 0 bytes, 0 discards
      0 CRC, 0 unknown class
      0 too long, 0 too short
    0 frames output, 0 bytes, 0 discards
    0 input OLS, 0 LRR, 0 NOS, 0 loop inits
    0 output OLS, 0 LRR, 0 NOS, 0 loop inits
```


Send documentation comments to mdsfeedback-doc@cisco.com.

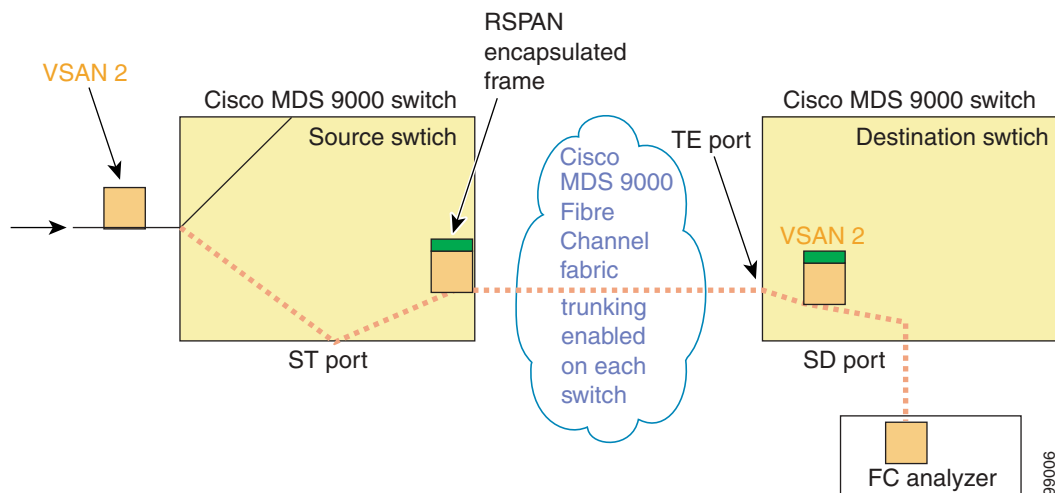
Remote SPAN

The Remote SPAN (RSPAN) feature enables you to remotely monitor traffic for one or more SPAN sources distributed in one or more source switches in a Fibre Channel fabric. The SPAN destination (SD) port is used for remote monitoring in a destination switch. A destination switch is usually different from the source switch(es) but is attached to the same Fibre Channel fabric. You can replicate and monitor traffic in any remote Cisco MDS 9000 Family switch or director, just as you would monitor traffic in a Cisco MDS source switch.

The RSPAN feature is nonintrusive and does not affect network traffic switching for that SPAN source ports. Traffic captured on the remote switch is tunneled across a Fibre Channel fabric which has trunking enabled on all switches in the path from the source switch to the destination switch. The Fibre Channel tunnel is structured using trunked ISL (TE) ports. In addition to TE ports, the RSPAN feature uses two other interface types (see [Figure 38-8](#)):

- SD ports—A passive port from which remote SPAN traffic can be obtained by the FC analyzer.
- ST ports—A SPAN tunnel (ST) port is an entry point port in the source switch for the RSPAN Fibre Channel tunnel. ST ports are special RSPAN ports and cannot be used for normal Fibre Channel traffic.

Figure 38-8 RSPAN Transmission



Advantages to Using RSPAN

The RSPAN features has the following advantages:

- Enables nondisruptive traffic monitoring at a remote location.
- Provides a cost effective solution by using one SD port to monitor remote traffic on multiple switches.
- Works with any Fibre Channel analyzer.
- Is compatible with the Cisco MDS 9000 Port Analyzer adapters.
- Does not affect traffic in the source switch, but shares the ISL bandwidth with other ports in the fabric.

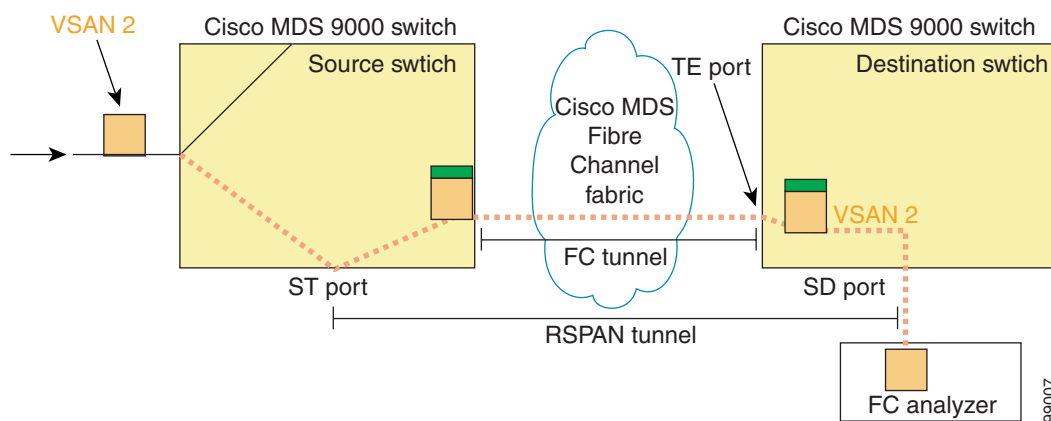
Send documentation comments to mdsfeedback-doc@cisco.com.

FC and RSPAN Tunnels

An FC tunnel is a logical data path between a source switch and a destination switch. The FC tunnel originates from the source switch and terminates at the remotely located destination switch.

RSPAN uses a special Fibre Channel tunnel (FC tunnel) that originates at the ST port in the source switch and terminates at the SD port in the destination switch. You must bind the FC tunnel to an ST port in the source switch and map the same FC tunnel to an SD port in the destination switch. Once the mapping and binding is configured, the FC tunnel is referred to as an RSPAN tunnel (see [Figure 38-9](#)).

Figure 38-9 FC and RSPAN Tunnel



Guidelines to Configure RSPAN

The following guidelines apply for a SPAN configuration:

- All switches in the end-to-end path of the RSPAN tunnel must belong to the Cisco MDS 9000 Family.
- All VSANs with RSPAN traffic must be enabled. If a VSAN containing RSPAN traffic is not enabled, it is dropped.
- The following configurations must be performed on *each* switch in the end-to-end path of the Fibre Channel tunnel in which RSPAN is to be implemented
 - Trunking must be enabled (enabled by default).
 - VSAN interface must be configured.
 - The Fibre Channel tunnel feature must be enabled (disabled by default).
 - IP routing must be enabled (disabled by default).



Note

If the IP address is in the same subnet as the VSAN, the VSAN interface does not have to be configured for all VSANs on which the traffic is spanned.

- A single Fibre Channel switch port must be dedicated for the ST port functionality.
- Do not configure the port to be monitored as the ST port.
- The FC tunnel's IP address must reside in the same subnet as the VSAN interface

See [Chapter 26, "Configuring IP Services."](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

ST Port Characteristics

ST port have the following characteristics:

- ST ports perform the RSPAN encapsulation of the FC frame.
- ST ports do not use BB_credits.
- One ST port can only be bound to one FC tunnel.
- ST ports cannot be used for any purpose other than to carry RSPAN traffic.
- ST Ports cannot be configured using Advanced Services Modules (ASMs) or Storage Services Modules (SSMs).

Configuring RSPAN

The RSPAN tunnel begins in the source switch and terminates in the destination switch. This section assumes Switch S to be the source and Switch D to be the destination.



Note

Besides the source and destination switches, the VSAN must also be configured in each Cisco MDS switch in the Fibre Channel fabric, if they exist.

To monitor network traffic using the RSPAN feature, follow these steps:

-
- | | |
|---------------|---|
| Step 1 | Create VSAN interfaces in destination switch (Switch D) and source switch (Switch S) to facilitate the Fibre Channel tunnel (FC tunnel) creation. |
| Step 2 | Enable the FC tunnel in each switch in the end-to-end path of the tunnel. |
| Step 3 | Initiate the FC tunnel (in Switch S) and map the tunnel to the VSAN interface's IP address (in Switch D) so all RSPAN traffic from the tunnel is directed to the SD port. |
| Step 4 | Configure SD ports for SPAN monitoring in the destination switch (Switch D). |
| Step 5 | Configure the ST port in the source switch (Switch S) and bind the ST port to the FC tunnel. |
| Step 6 | Create an RSPAN session in the source switch (in Switch S) to monitor network traffic. |
-

RSPAN Configuration Example

This section provides a RSPAN configuration example using the procedure defined in the previous section.

Configuration in the Source Switch


This section identifies the tasks that must be performed in the source switch (Switch D):

Send documentation comments to mdsfeedback-doc@cisco.com.

Creating VSAN Interfaces

Figure 38-10 depicts a basic FC tunnel configuration.

Figure 38-10 FC Tunnel Configuration



Note

This example assumes that VSAN 5 is already configured in the VSAN database.

To create a VSAN interface in the source switch for the scenario in Figure 38-10, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switchS# config t | Enters configuration mode. |
| Step 2 | switchS(config)# interface vsan 5 switchS(config-if)# | Configures the specified VSAN interface (VSAN 5) in the source switch (switch S). |
| Step 3 | switchS(config-if)# ip address 10.10.10.1 255.255.255.0 | Configures the IP address and subnet for the VSAN interface 5 in the source switch (switch S). |
| Step 4 | switchS(config-if)# no shutdown | Enables traffic flow through this interface. |

Enabling FC Tunnels

To enable the FC tunnel feature, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switchS# config t | Enters configuration mode. |
| Step 2 | switchS(config)# fc-tunnel enable | Enables the FC tunnel feature (disabled by default). |



Note

Be sure to enable this feature in each switch in the end-to-end path in the fabric.

Initiating the FC Tunnel

To initiate the FC tunnel in the source switch for the scenario in Figure 38-10, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switchS# config t | Enters configuration mode. |
| Step 2 | switchS(config)# interface fc-tunnel 100 switchS(config-if)# | Initiates the FC tunnel (100) in the source switch (switch S). The tunnel IDs range from 1 to 255. |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | Command | Purpose |
|--------|---|--|
| Step 3 | switchS(config-if)# source 10.10.10.1 | Maps the IP address of the source switch (switch S) to the FC tunnel (100). |
| Step 4 | switchS(config-if)# destination 10.10.10.2 | Maps the IP address of the destination switch (switch D) to the FC tunnel (100). |
| Step 5 | switchS(config-if)# no shutdown | Enables traffic flow through this interface. |



Tip


The interface cannot be operationally up until the FC tunnel mapping is configured in the destination switch.

Configuring the ST Port

Once the FC tunnel is created, be sure to configure the ST port to bind it to the FC tunnel at the source switch. The FC tunnel becomes an RSPAN tunnel once the binding and mapping is complete.

Figure 38-11 depicts a basic FC tunnel configuration.

Figure 38-11 Binding the FC Tunnel



Note

ST ports cannot be configured using Advanced Services Modules (ASMs) or Storage Services Modules (SSMs).

To configure an ST port for the scenario in Figure 38-11, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switchS# config t | Enters configuration mode. |
| Step 2 | switchS(config)# interface fc2/1 | Configures the specified interface. |
| Step 3 | switchS(config-if)# switchport mode ST | Configures the ST port mode for interface fc2/1. |
| Step 4 | switchS(config-if)# switchport speed 2000 | Configures the ST port speed to 2000 Mbps. |
| Step 5 | switchS(config-if)# rspan-tunnel interface fc-tunnel 100 | Associates and binds the ST port with the RSPAN tunnel (100). |
| Step 6 | switchS(config-if)# no shutdown | Enables traffic flow through this interface. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Configure an RSPAN Session

A RSPAN session is similar to a SPAN session, with the destination interface being an RSPAN tunnel. To configure an RSPAN session in the source switch for the scenario in [Figure 38-11](#), follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switchS# config t | Enters configuration mode. |
| Step 2 | switchS(config)# span session 2 switchS(config-span)# | Configures the specified SPAN session (2). If the session does not exist, it is created. The session ID ranges from 1 to 16. |
| Step 3 | switchS(config-span)# destination interface fc-tunnel 100 | Configures the specified RSPAN tunnel (100) in a session. |
| Step 4 | switchS(config-span)# source interface fc1/1 | Configures the source interface (fc1/1) for this session and spans the traffic from interface fc1/1 to RSPAN tunnel 100. |

Configuration in All Intermediate Switches

This section identifies the tasks that must be performed in all intermediate switches in the end-to-end path of the RSPAN tunnel:

- [Configuring VSAN Interfaces, page 38-20](#)
- [Enabling FC Tunnels, page 38-21](#)
- [Enabling IP Routing, page 38-21](#)

Configuring VSAN Interfaces

[Figure 38-13](#) depicts an RSPAN tunnel configuration terminating in the destination switch (Switch D).



Note

This example assumes that VSAN 5 is already configured in the VSAN database.

To create a VSAN interface in the destination switch for the scenario in [Figure 38-13](#), follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switchD# config t | Enters configuration mode. |
| Step 2 | switchD(config)# interface vsan 5 switchD(config-if)# | Configures the specified VSAN interface (VSAN 5) in the destination switch (Switch D). |
| Step 3 | switchD(config-if)# ip address 10.10.10.2 255.255.255.0 | Configures the IP address and subnet for the VSAN interface in the destination switch (Switch D). |
| Step 4 | switchD(config-if)# no shutdown | Enables traffic flow to administratively allow traffic (provided the operational state is up). |

Send documentation comments to mdsfeedback-doc@cisco.com.

Enabling FC Tunnels

To enable the FC tunnel feature, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switchS# config t | Enters configuration mode. |
| Step 2 | switchS(config)# fc-tunnel enable | Initiates the FC tunnel (100) in the source switch (switch S). The tunnel IDs range from 1 to 255. |



Note

Be sure to enable this feature in each switch in the end-to-end path in the fabric.

Enabling IP Routing

The IP routing feature is disabled by default. Be sure to enable IP routing in each switch (including the source and destination switches) in the end-to-end path in the fabric (see [“Enabling IP Routing” section on page 26-12](#)). This step is required to set up the FC tunnel.

Configuration in the Destination Switch

This section identifies the tasks that must be performed in the destination switch (Switch D):

- [Configuring VSAN Interfaces, page 38-21](#)
- [Configuring the SD Port, page 38-22](#)
- [Mapping the FC Tunnel, page 38-23](#)

Configuring VSAN Interfaces

[Figure 38-12](#) depicts an RSPAN tunnel configuration terminating in the destination switch (Switch D).



Note

This example assumes that VSAN 5 is already configured in the VSAN database.

To create a VSAN interface in the destination switch for the scenario in [Figure 38-12](#), follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switchD# config t | Enters configuration mode. |
| Step 2 | switchD(config)# interface vsan 5 switchD(config-if)# | Configures the specified VSAN interface (VSAN 5) in the destination switch (Switch D). |
| Step 3 | switchD(config-if)# ip address 10.10.10.2 255.255.255.0 | Configures the IP address and subnet for the VSAN interface in the destination switch (Switch D). |
| Step 4 | switchD(config-if)# no shutdown | Enables traffic flow to administratively allow traffic (provided the operational state is up). |

Send documentation comments to mdsfeedback-doc@cisco.com.

Enabling FC Tunnels

To enable the FC tunnel feature, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switchS# config t | Enters configuration mode. |
| Step 2 | switchS(config)# fc-tunnel enable | Initiates the FC tunnel (100) in the source switch (switch S). The tunnel IDs range from 1 to 255. |



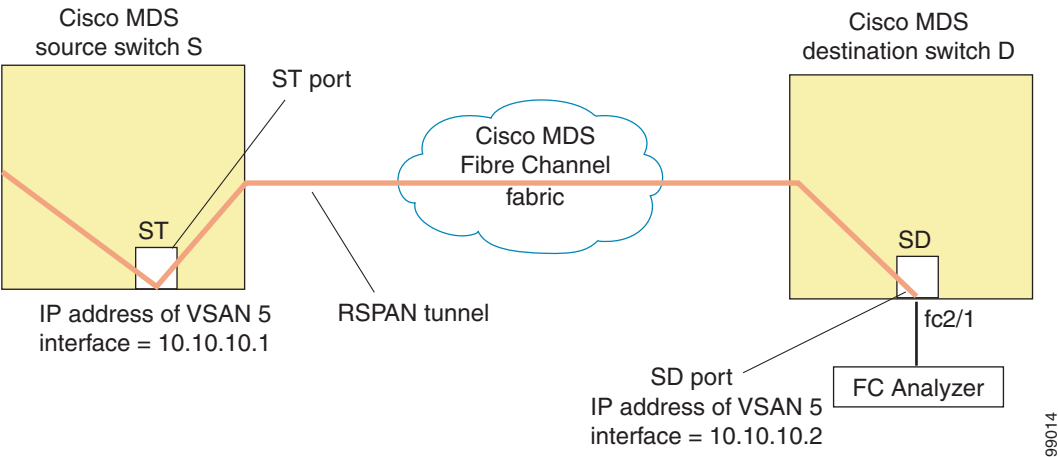
Note

Be sure to enable this feature in each switch in the end-to-end path in the tunnel.

Configuring the SD Port

The SD port in the destination switch enables the FC analyzer to receive the RSPAN traffic from the Fibre Channel tunnel. Figure 38-12 depicts an RSPAN tunnel configuration, now that tunnel destination is also configured.

Figure 38-12 RSPAN Tunnel Configuration



Note

SD ports cannot be configured using Advanced Services Modules (ASMs) or Storage Services Modules (SSMs).

To configure an SD port for the scenario in Figure 38-12, follow these steps:

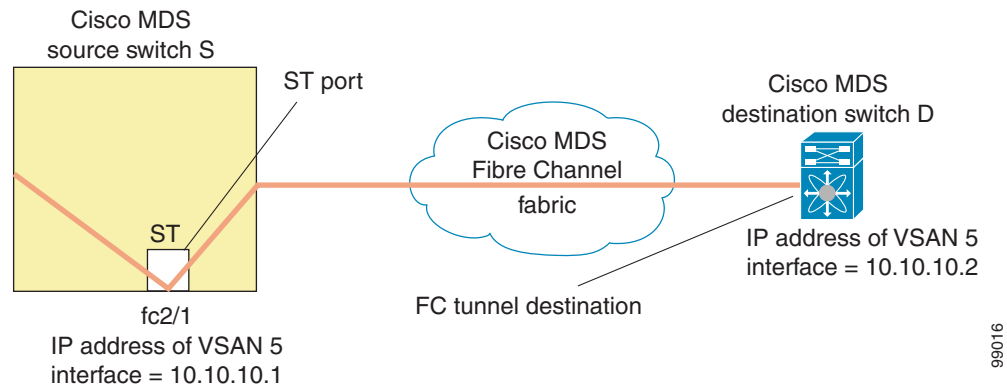
| | Command | Purpose |
|--------|--|--|
| Step 1 | switchD# config t | Enters configuration mode. |
| Step 2 | switchD(config)# interface fc2/1 | Configures the specified interface. |
| Step 3 | switchD(config-if)# switchport mode SD | Configures the SD port mode for interface fc2/1. |
| Step 4 | switchD(config-if)# switchport speed 2000 | Configures the SD port speed to 2000 Mbps. |
| Step 5 | switchD(config-if)# no shutdown | Enables traffic flow through this interface. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Mapping the FC Tunnel

The **tunnel-id-map** option specifies the egress interface of the tunnel at the destination switch (see Figure 38-13).

Figure 38-13 FC Tunnel Configuration



To terminate the FC tunnel in the destination switch for the scenario in Figure 38-13, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switchD# conf t | Enters configuration mode. |
| Step 2 | switchD(config)# fc-tunnel tunnel-id-map 100 interface fc2/1 | Terminates the FC tunnel (100) in the destination switch (switch D). The tunnel ID range is from 1 to 255. |


Explicit Paths

You can specify an explicit path through the Cisco MDS Fibre channel fabric (source-based routing), using the **explicit-path** option. For example, if you have multiple paths to a tunnel destination, you can use this option to specify the fc-tunnel to always take one path to the destination switch. The software then uses this specified path even if other paths are available.

This option is especially useful if you prefer to direct the traffic through a certain path although other paths are available. In an RSPAN situation, you can specify the explicit path so the RSPAN traffic does not interfere with the existing user traffic. You can create any number of explicit paths in a switch (see Figure 38-14).

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 38-14 Explicit Path Configuration



The explicit path must be created in the source switch. To configure an explicit path, you must first create the path and then configure the use of any one path. If an explicit path is not configured, the minimum cost path is used by default. If an explicit path is configured and is functioning, the specified path is used.

To create an explicit path for the scenario in [Figure 38-14](#), follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switchS# config t | Enters configuration mode. |
| Step 2 | switchS(config)# fc-tunnel explicit-path Path1 switch(config-explicit-path)# | Places you at the explicit path prompt for the path named Path 1. |
| Step 3 | switchS(config-explicit-path)# next-address 10.10.10.2 strict switchS(config-explicit-path)# next-address 10.10.10.3 strict switchS(config-explicit-path)# next-address 10.10.10.4 strict | Specifies that the next hop VSAN interface IP addresses and the previous hops specified in the explicit path do not require direct connection. |
| Step 4 | switchS(config)# fc-tunnel explicit-path Path2 switch(config-explicit-path)# | Places you at the explicit path prompt for Path2. |
| Step 5 | switchS(config-explicit-path)# next-address 10.10.10.5 strict switchS(config-explicit-path)# next-address 10.10.10.4 strict | Specifies that the next hop VSAN interface IP addresses and the previous hops specified in the explicit path does not require direct connection. |
| Step 6 | switchS(config)# fc-tunnel explicit-path Path3 switch(config-explicit-path)# | Places you at the explicit path prompt for Path3. |
| Step 7 | switchS(config-explicit-path)# next-address 10.10.10.3 loose | Configures a minimum cost path in which the 10.10.10.3 IP address exists. Note In Figure 38-14 , Path3 is the same as Path1—10.10.10.3 exists in Path 1. Using the loose option, you can achieve the same effect with one command instead of issuing three commands (using the strict option) in Step 3. |

Send documentation comments to mdsfeedback-doc@cisco.com.

To reference the explicit path, follow these steps:

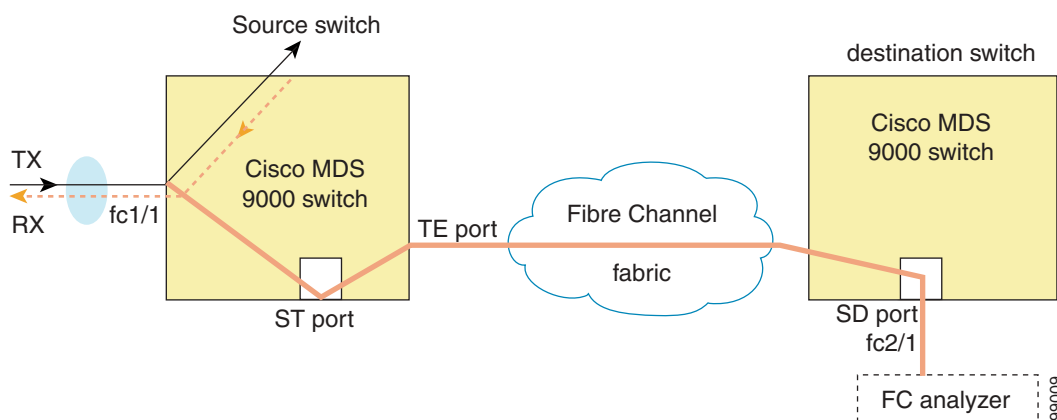
| | Command | Purpose |
|--------|---|-------------------------------------|
| Step 1 | switchS# config t | Enters configuration mode. |
| Step 2 | switchS(config)# interface fc-tunnel 100 | References the tunnel ID for Path1. |
| Step 3 | switchS(config)# explicit-path Path1 | Links Path1 to the tunnel ID. |

This configuration explicitly specifies Path 1 to be used for the RSPAN traffic. Refer to RFC 3209 for further details on explicit paths and source based routing.

Monitoring RSPAN Traffic

Once the session is configured, other SPAN sources for this session can also be configured as required. [Figure 38-15](#) shows an RSPAN setup where one session with destination port fc2/1 and source interface fc1/1 is used to capture traffic in both ingress and egress directions.

Figure 38-15 Fibre Channel Analyzer Using a Single SD Port to Monitor RSPAN Traffic



To use this setup, the analyzer should have the capability of distinguishing ingress and egress traffic for all captured frames.

Sample Scenarios



Note


RSPAN can be combined with the local SPAN feature so SD ports forward local SPAN traffic along with remote SPAN traffic. Various SPAN source and tunnel scenarios are described in this section.

Send documentation comments to mdsfeedback-doc@cisco.com.

Single Source with One RSPAN Tunnel

The source Switch S and the destination Switch D are interconnected through a Fibre Channel fabric. A RSPAN tunnel is configured as a destination interface for the SPAN session and the ST port forwards SPAN traffic through the RSPAN tunnel (see Figure 38-16).


Figure 38-16 RSPAN Scenario with One Source Switch, One Destination Switch, and One Tunnel



Single Source with Multiple RSPAN Tunnels

Figure 38-17 displays two separate RSPAN tunnels configured between Switches S and N. Each tunnel has an associated ST port in the source switch and a separate SD port in the destination switch. This configuration is useful for troubleshooting purposes.

Figure 38-17 RSPAN Scenario with One Source Switch, One Destination Switch, and Multiple Tunnels

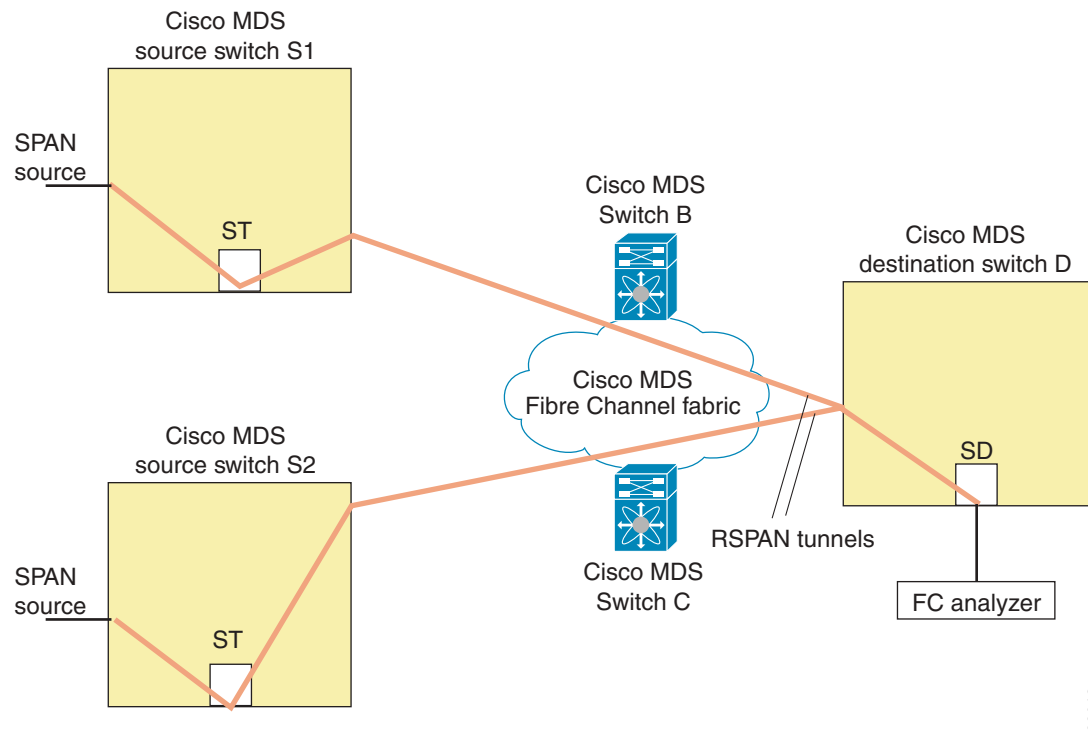


Send documentation comments to mdsfeedback-doc@cisco.com.

Multiple Sources with Multiple RSPAN Tunnels

Figure 38-18 displays two separate RSPAN tunnels configured between Switches S1 and S2. Both tunnels have an associated ST port in their respective source switch and terminate in the same SD port in the destination switch.

Figure 38-18 RSPAN Scenario with Two Source Switches, a Destination Switch, and Multiple Tunnels



This configuration is useful for remote monitoring purposes. For example, the administrator may be at the destination switch and can remotely monitor the two source switches.

Displaying RSPAN Information

Use the **show** commands to display configured RSPAN information. See Examples 38-5 to 38-11.

Example 38-5 Displays ST Port Interface Information

```
switch# show interface brief
```

| Interface | Vsan | Admin Mode | Admin Trunk Mode | Status | Oper Mode | Oper Speed (Gbps) | Port-channel |
|---------------|----------|------------|------------------|-----------|-----------|-------------------|--------------|
| fc1/1 | 1 | auto | on | trunking | TE | 2 | -- |
| ... | | | | | | | |
| fc1/14 | 1 | auto | on | trunking | TE | 2 | -- |
| fc1/15 | 1 | ST | on | up | ST | 2 | -- |
| ... | | | | | | | |

Send documentation comments to mdsfeedback-doc@cisco.com.

| | | | | | | | |
|-----------------|-----|--------|------------------|-------------|--------------|-------------------|-----------------|
| fc2/9 | 1 | auto | on | trunking | TE | 2 | port-channel 21 |
| fc2/10 | 1 | auto | on | trunking | TE | 2 | port-channel 21 |
| ... | | | | | | | |
| fc2/13 | 999 | auto | on | up | F | 1 | -- |
| fc2/14 | 999 | auto | on | up | FL | 1 | -- |
| fc2/15 | 1 | SD | -- | up | SD | 2 | -- |
| fc2/16 | 1 | auto | on | trunking | TE | 2 | -- |
| ----- | | | | | | | |
| Interface | | Status | | | Speed (Gbps) | | |
| ----- | | | | | | | |
| sup-fc0 | | up | | | 1 | | |
| ----- | | | | | | | |
| Interface | | Status | IP Address | | Speed | MTU | |
| ----- | | | | | | | |
| mgmt0 | | up | 172.22.36.175/22 | | 100 Mbps | 1500 | |
| ----- | | | | | | | |
| Interface | | Status | IP Address | | Speed | MTU-- | |
| ----- | | | | | | | |
| vsan5 | | up | 10.10.10.1/24 | | 1 Gbps | 1500 | |
| ----- | | | | | | | |
| Interface | | Vsan | Admin Trunk Mode | Status | Oper Mode | Oper Speed (Gbps) | |
| ----- | | | | | | | |
| port-channel 21 | | 1 | on | trunking | TE | 4 | |
| ----- | | | | | | | |
| Interface | | Status | Dest IP Addr | Src IP Addr | TID | Explicit Path | |
| ----- | | | | | | | |
| fc-tunnel 100 | | up | 10.10.10.2 | 10.10.10.1 | 100 | | |

Example 38-6 Displays Detailed Information for the ST Port Interface

```
switch# show interface fc1/11
fc1/11 is up
  Hardware is Fibre Channel
  Port WWN is 20:0b:00:05:30:00:59:de
  Admin port mode is ST
  Port mode is ST
  Port vsan is 1
  Speed is 1 Gbps
  Rspan tunnel is fc-tunnel 100
  Beacon is turned off
  5 minutes input rate 248 bits/sec, 31 bytes/sec, 0 frames/sec
  5 minutes output rate 176 bits/sec, 22 bytes/sec, 0 frames/sec
  6862 frames input, 444232 bytes
    0 discards, 0 errors
    0 CRC, 0 unknown class
    0 too long, 0 too short
  6862 frames output, 307072 bytes
    0 discards, 0 errors
  0 input OLS, 0 LRR, 0 NOS, 0 loop inits
  0 output OLS, 0 LRR, 0 NOS, 0 loop inits
```

Example 38-7 Displays the FC Tunnel Status

```
switch# show fc-tunnel
fc-tunnel is enabled
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 38-8 Displays FC Tunnel Egress Mapping Information

```
switch# show fc-tunnel tunnel-id-map
tunnel id egress interface
      150      fc3/1
      100      fc3/1
```



Note

Multiple tunnel IDs can terminate at the same interface.

Example 38-9 Displays FC Tunnel Explicit Mapping Information

```
switch# show fc-tunnel explicit-path
Explicit path name: Alternatel
      10.20.1.2 loose
      10.20.1.3 strict
Explicit path name: User2
      10.20.50.1 strict
      10.20.50.4 loose
```

Example 38-10 Displays SPAN Mapping Information

```
switch# show span session
Session 2 (active)
  Destination is fc-tunnel 100
  No session filters configured
  Ingress (rx) sources are
    fc2/16,
  Egress (tx) sources are
    fc2/16,
```

Example 38-11 Displays the FC Tunnel Interface

```
switch# show interface fc-tunnel 200
fc-tunnel 200 is up
Dest   IP Addr: 200.200.200.7   Tunnel ID: 200
Source IP Addr: 200.200.200.4 LSP ID: 1
Explicit Path Name:
```

Default SPAN and RSPAN Settings

Table 38-1 lists the default settings for SPAN parameters.

Table 38-1 Default SPAN Configuration Parameters

| Parameters | Default |
|------------------------------|---|
| SPAN session | Active. |
| If filters are not specified | SPAN traffic includes traffic through a specific interface from all active VSANs. |
| Encapsulation | Disabled. |
| SD port | Output frame format is Fibre Channel. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 38-2 lists the default settings for RSPAN parameters.

Table 38-2 ***Default SPAN Configuration Parameters***

| Parameters | Default |
|-------------------|--|
| FC tunnel | Disabled. |
| Explicit path | Not configured. |
| Minimum cost path | Used if explicit path is not configured. |

Default RSPAN Settings

Table 38-3 lists the default settings for RSPAN parameters.

Table 38-3 ***Default RSPAN Configuration Parameters***

| Parameters | Default |
|-------------------|--|
| FC tunnel | Disabled. |
| Explicit path | Not configured. |
| Minimum cost path | Used if explicit path is not configured. |



Advanced Features and Concepts

This chapter describes the advanced features provided in switches in the Cisco MDS 9000 Family. It includes the following sections:

- [Fibre Channel Time Out Values, page 39-2](#)
- [The fctrace Feature, page 39-5](#)
- [The fcping Feature, page 39-7](#)
- [Configuring a Fabric Analyzer, page 39-8](#)
- [Configuring World Wide Names, page 39-18](#)
- [Configuring a Secondary MAC Address, page 39-19](#)
- [FC ID Allocation for HBAs, page 39-20](#)
- [Loop Monitoring Initiation, page 39-22](#)
- [Switch Interoperability, page 39-22](#)
- [The show tech-support Command, page 39-29](#)
- [Default Settings, page 39-31](#)

[Send documentation comments to mdsfeedback-doc@cisco.com.](mailto:mdsfeedback-doc@cisco.com)

Fibre Channel Time Out Values

You can modify Fibre Channel protocol related timer values for the switch by configuring the following time out values (TOVs):

- Distributed services TOV (D_S_TOV)—The valid range is from 5,000 to 10,000 milliseconds. The default is 5,000 milliseconds.
- Error detect TOV (E_D_TOV)—The valid range is from 1,000 to 10,000 milliseconds. The default is 2,000 milliseconds. This value is matched with the other end during port initialization.
- Resource allocation TOV (R_A_TOV)—The valid range is from 5,000 to 10,000 milliseconds. The default is 10,000 milliseconds. This value is matched with the other end during port initialization.



Note

The fabric stability TOV (F_S_TOV) constant cannot be configured.

Timer Configuration Across All VSANs

You can modify Fibre Channel protocol related timer values for the switch.



Caution

The D_S_TOV, E_D_TOV, and R_A_TOV values cannot be globally changed unless all VSANs in the switch are suspended.



Note

If a VSAN is not specified when you change the timer value, the changed value is applied to all VSANs in the switch.

To configure FC timers across all VSANs, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config) | Enters configuration mode. |
| Step 2 | switch(config)# fctimer R_A_TOV 6000 | Configures the R_A_TOV value for all VSANs to be 6000 ms. This type of configuration is not permitted unless all VSANs are suspended. |

Timer Configuration Per-VSAN

You can also issue the fctimer for a specified VSAN to configure different TOV values for VSANs with special links like FC or IP tunnels. You can configure different E_D_TOV, R_A_TOV, and D_S_TOV values for individual VSANs. Active VSANs are suspended and activated when their timer values are changed.



Caution

You cannot perform a nondisruptive downgrade to any earlier version that does not support per-VSAN FC timers.

Send documentation comments to mdsfeedback-doc@cisco.com.



Note

This configuration must be propagated to all switches in the fabric—be sure to configure the same value in all switches in the fabric.

If a switch is downgraded to Cisco MDS SAN-OS Release 1.2 or 1.1 after the timer is configured for a VSAN, an error message is issued to warn against strict incompatibilities. Refer to the *Cisco MDS 9000 Family Troubleshooting Guide*.

To configure per-VSAN FC timers, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config) | Enters configuration mode. |
| Step 2 | switch(config)# fctimer D_S_TOV 6000 vsan 2 Warning: The vsan will be temporarily suspended when updating the timer value This configuration would impact whole fabric. Do you want to continue? (y/n) y Since this configuration is not propagated to other switches, please configure the same value in all the switches | Configures the D_S_TOV value to be 6000 ms for VSAN 2. Suspends the VSAN temporarily. You have the option to end this command, if required. |

fctimer Distribution

As of Cisco SAN-OS Release 2.0(1b), you can enable per-VSAN fctimer fabric distribution for all Cisco MDS switches in the fabric. When you perform fctimer configurations, and distribution is enabled, that configuration is distributed to all the switches in the fabric.

You automatically acquire a fabric-wide lock when you issue the first configuration command after you enabled distribution in a switch. The fctimer application uses the effective and pending database model to store or commit the commands based on your configuration.

Refer to [Chapter 9, “Using the CFS Infrastructure”](#) for more information on the CFS application.

To enable fctimer fabric distribution, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# fctimer distribute | Enables fctimer configuration distribution to all switches in the fabric. Acquires a fabric lock and stores all future configuration changes in the pending database. |
| | switch(config)# no fctimer distribute | Disables (default) fctimer configuration distribution to all switches in the fabric. |

Committing fctimer Changes

When you commit the fctimer configuration changes, the effective database is overwritten by the configuration changes in the pending database and all the switches in the fabric receive the same configuration. When you commit the fctimer configuration changes without implementing the session feature, the fctimer configurations are distributed to all the switches in the physical fabric.

Send documentation comments to mdsfeedback-doc@cisco.com.

To commit the fctimer configuration changes, follow these steps:

| | Command | Purpose |
|--------|---------------------------------------|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# fctimer commit | Distributes the fctimer configuration changes to all switches in the fabric and releases the lock. Overwrites the effective database with the changes made to the pending database. |

Discarding fctimer Changes

After making the configuration changes, you can choose to discard the changes by aborting the changes instead of committing them. In either case, the lock is released.

To discard the fctimer configuration changes, follow these steps:

| | Command | Purpose |
|--------|--------------------------------------|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# fctimer abort | Discards the fctimer configuration changes in the pending database and releases the fabric lock. |

Fabric Lock Override

If you have performed a fctimer fabric task and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.



Tip

The changes are only available in the volatile directory and are subject to being discarded if the switch is restarted.

To use administrative privileges and release a locked fctimer session, use the **clear fctimer session** command.

```
switch# clear fctimer session
```

Database Merge Guidelines

See the “CFS Merge Support” section on page 9-7 for detailed concepts.

Send documentation comments to mdsfeedback-doc@cisco.com.

When merging two fabrics, follow these guidelines:

- Be aware of the following merge conditions:
 - The merge protocol is not implemented for distribution of the fctimer values—you must manually merge the fctimer values when a fabric is merged. The per-VSAN fctimer configuration is distributed in the physical fabric.
 - The fctimer configuration is only applied to those switches containing the VSAN with a modified fctimer value.
 - The global fctimer values are not distributed.
- Do not configure global timer values when distribution is enabled.



Note

The number of pending fctimer configuration operations cannot be more than 15. At that point, you must commit or abort the pending configurations before performing any more operations.

Displaying Configured FC Timer Values

Use the **show fctimer** command to display the configured FC timer values (see Examples 39-1 and 39-2).

Example 39-1 Displays Configured Global TOVs

```
switch# show fctimer
F_S_TOV   D_S_TOV   E_D_TOV   R_A_TOV
-----
5000 ms   5000 ms   2000 ms   10000 ms
```



Note

The F_S_TOV constant, though not configured, is displayed in the output of the **show fctimer** command.

Example 39-2 Displays Configured TOVs for a Specified VSAN

```
switch# show fctimer vsan 10
vsan no.  F_S_TOV   D_S_TOV   E_D_TOV   R_A_TOV
-----
10         5000 ms   5000 ms   3000 ms   10000 ms
```

The fctrace Feature

The fctrace feature allows you to:

- Trace the route followed by data traffic.
- Compute inter-switch (hop-to-hop) latency.

You can invoke fctrace by providing the FC ID, the N port, or the NL port WWN, or the device alias of the destination. The frames are routed normally as long as they are forwarded through TE ports.

Once the frame reaches the edge of the fabric (the F port or FL port connected to the end node with the given port WWN or the FC ID), the frame is looped back (swapping the source ID and the destination ID) to the originator.

If the destination cannot be reached, the path discovery starts, which traces the path up to the point of failure.

Send documentation comments to mdsfeedback-doc@cisco.com.

**Note**

The fctrace feature works only on TE ports. Make sure that only TE ports exist in the path to the destination. In case there is an E port in the path, the fctrace frame is dropped by that switch. Also, fctrace times out in the originator, and path discovery does not start.

**Tip**

You cannot use the fctrace feature in a locally configured VSAN interface (IPFC interface), but you can trace the route to a VSAN interface configured in other switches.

To perform a fctrace operation, follow this step:

| | Command | Purpose |
|---------------|---|--|
| Step 1 | <pre>switch# fctrace fcid 0xd70000 vsan 1 Route present for : 0xd70000 20:00:00:0b:46:00:02:82(0xffffcd5) Timestamp Invalid. 20:00:00:05:30:00:18:db(0xffffcd7) Timestamp Invalid. 20:00:00:05:30:00:18:db(0xffffcd7)</pre> | Invokes fctrace for the specified FC ID of the destination N port. |
| | <pre>switch# fctrace pwwn 21:00:00:e0:8b:06:d9:1d vsan 1 timeout 5 Route present for : 21:00:00:e0:8b:06:d9:1d 20:00:00:0b:46:00:02:82(0xffffcd5) Timestamp Invalid. 20:00:00:05:30:00:18:db(0xffffcd7) Timestamp Invalid. 20:00:00:05:30:00:18:db(0xffffcd7)</pre> | Invokes fctrace using the pWWN of the destination N port. By default the period to wait before timing out is 5 seconds. The range is from one through 10 seconds. |
| | <pre>switch# fctrace device-alias disk1 v 1 Route present for : 22:00:00:0c:50:02:ce:f8 20:00:00:05:30:00:31:1e(0xffffca9)</pre> | Invokes fctrace using the device alias of the destination N port. |

Send documentation comments to mdsfeedback-doc@cisco.com.

The fcping Feature

The fcping feature verifies reachability of a node by checking its end-to-end connectivity. You can invoke the fcping feature by providing the FC ID, the destination port WWN, or the device alias information.

To perform a fcping operation, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | <pre>switch# fcping fcid 0xd70000 vsan 1 28 bytes from 0xd70000 time = 730 usec 28 bytes from 0xd70000 time = 165 usec 28 bytes from 0xd70000 time = 262 usec 28 bytes from 0xd70000 time = 219 usec 28 bytes from 0xd70000 time = 228 usec 5 frames sent, 5 frames received, 0 timeouts Round-trip min/avg/max = 165/270/730 usec</pre> | Invokes fcping for the specified pWWN or the FC ID of the destination. By default, five frames are sent. |
| | <pre>switch# fcping fcid 0xd70000 vsan 1 count 10 28 bytes from 0xd70000 time = 730 usec 28 bytes from 0xd70000 time = 165 usec 28 bytes from 0xd70000 time = 262 usec 28 bytes from 0xd70000 time = 219 usec 28 bytes from 0xd70000 time = 228 usec 28 bytes from 0xd70000 time = 230 usec 28 bytes from 0xd70000 time = 230 usec 28 bytes from 0xd70000 time = 225 usec 28 bytes from 0xd70000 time = 229 usec 28 bytes from 0xd70000 time = 183 usec 10 frames sent, 10 frames received, 0 timeouts Round-trip min/avg/max = 165/270/730 usec</pre> | Sets the number of frames to be sent using the count option. The range is from 0 through 2147483647. A value of 0 pings forever. |
| | <pre>switch# fcping fcid 0xd500b4 vsan 1 timeout 10 28 bytes from 0xd500b4 time = 1345 usec ... 5 frames sent, 5 frames received, 0 timeouts Round-trip min/avg/max = 340/581/1345 usec</pre> | Sets the timeout value. The default period to wait is 5 seconds. The range is from 1 through 10 seconds. |
| | <pre>switch# fcping device-alias disk1 vsan 1 28 bytes from 22:00:00:0c:50:02:ce:f8 time = 1883 usec 28 bytes from 22:00:00:0c:50:02:ce:f8 time = 493 usec 28 bytes from 22:00:00:0c:50:02:ce:f8 time = 277 usec 28 bytes from 22:00:00:0c:50:02:ce:f8 time = 391 usec 28 bytes from 22:00:00:0c:50:02:ce:f8 time = 319 usec 5 frames sent, 5 frames received, 0 timeouts Round-trip min/avg/max = 277/672/1883 usec</pre> | Invokes fcping for the specified device alias of the destination. |
| Step 2 | <pre>switch# fcping fcid 0x010203 vsan 1 No response from the N port. switch# fcping pwn 21:00:00:20:37:6f:db:dd vsan 1 28 bytes from 21:00:00:20:37:6f:db:dd time = 1454 usec ... 5 frames sent, 5 frames received, 0 timeouts Round-trip min/avg/max = 364/784/1454 usec</pre> | <p>Issues a No response from the N port message even when the N port or NL port is active. This is due to resource exhaustion at the N port or NL port.</p> <p>Retry the command a few seconds later.</p> |

Send documentation comments to mdsfeedback-doc@cisco.com.

Verifying Switch Connectivity

You can verify connectivity to a destination switch.



Note

The FC ID variable used in this procedure is the domain controller address; it is not a duplication of the domain ID.

To verify connectivity to a destination switch, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | <pre>switch# show fcdomain domain-list vsan 200 Number of domains: 7 Domain ID WWN ----- 0x01(1) 20:c8:00:05:30:00:59:df [Principal] 0x02(2) 20:c8:00:0b:5f:d5:9f:c1 0x6f(111) 20:c8:00:05:30:00:60:df 0xda(218) 20:c8:00:05:30:00:87:9f [Local] 0x06(6) 20:c8:00:0b:46:79:f2:41 0x04(4) 20:c8:00:05:30:00:86:5f 0x6a(106) 20:c8:00:05:30:00:f8:e3</pre> | <p>Displays the destination switch's domain ID.</p> <p>To obtain the domain controller address, concatenate the domain ID with FFC. For example, if the domain ID is 0xda(218), the concatenated ID is 0xfffcda.</p> |
| Step 2 | <pre>switch# fcping fcid 0xFFFCDA vsan 200 28 bytes from 0xFFFCDA time = 298 usec 28 bytes from 0xFFFCDA time = 260 usec 28 bytes from 0xFFFCDA time = 298 usec 28 bytes from 0xFFFCDA time = 294 usec 28 bytes from 0xFFFCDA time = 292 usec 5 frames sent, 5 frames received, 0 timeouts Round-trip min/avg/max = 260/288/298 usec</pre> | <p>Verifies reachability of the destination switch by checking its end-to-end connectivity.</p> |

Configuring a Fabric Analyzer

Fibre Channel protocol analyzers capture, decode, and analyze frames and ordered sets on a link. Existing Fibre Channel analyzers can capture traffic at wire rate speed. They are expensive and support limited frame decoding. Also, to snoop traffic, the existing analyzers disrupt the traffic on the link while the analyzer is inserted into the link.

Cisco has brought protocol analysis within a storage network to a new level with the Cisco Fabric Analyzer. You can capture Fibre Channel control traffic from a switch and decode it without having to disrupt any connectivity, and without having to be local to the point of analysis.

The Cisco Fibre Channel protocol analyzer is based on two popular public-domain software applications:

- libpcap—See <http://www.tcpdump.org>.
- Ethereal—See <http://www.ethereal.com>.



Note

The Cisco Fabric Analyzer is useful in capturing and decoding control traffic, not data traffic. It is suitable for control path captures, and is not intended for high-speed data path captures.

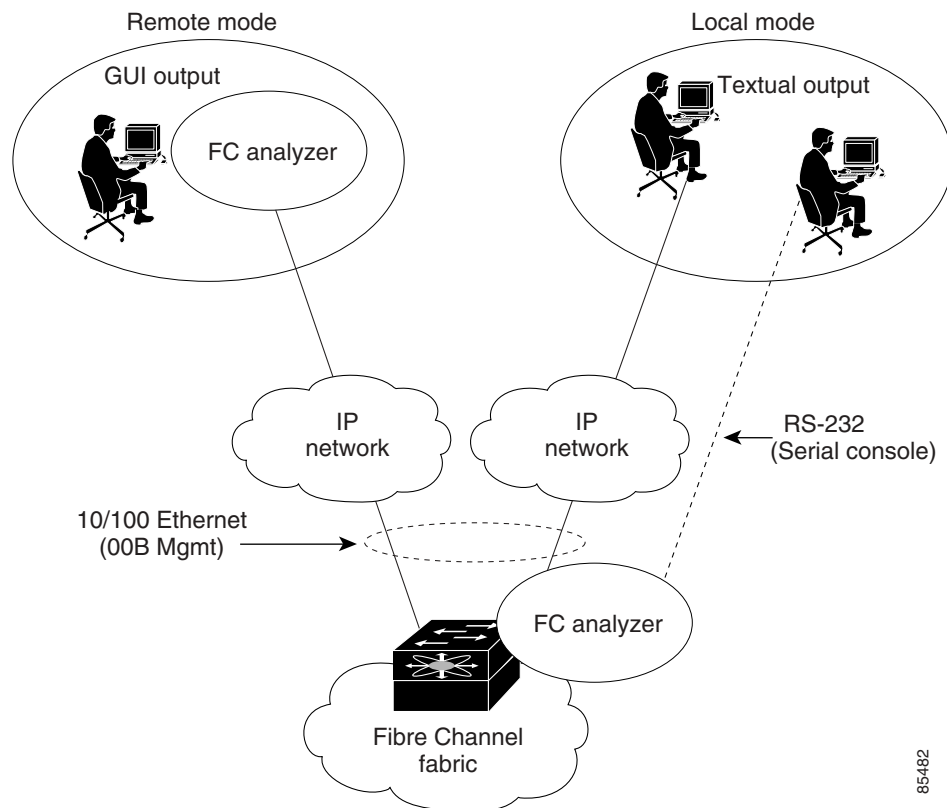
Send documentation comments to mdsfeedback-doc@cisco.com.

About the Cisco Fabric Analyzer

The Cisco Fabric Analyzer consists of two separate components (see [Figure 39-1](#)):

- Software that runs on the Cisco MDS 9000 Family switch and supports two modes of capture:
 - A text-based analyzer that supports local capture and decodes captured frames
 - A daemon that supports remote capture
- GUI-based client that runs on a host that supports libpcap such as Windows or Linux and communicates with the remote capture daemon in a Cisco MDS 9000 Family switch.

Figure 39-1 Cisco Fabric Analyzer Usage



Local Text-Based Capture

This component is a command-line driven text-based interface that captures traffic to and from the supervisor module in a Cisco MDS 9000 Family switch. It is a fully functional decoder that is useful for quick debug purposes or for use when the remote capture daemon is not enabled. Additionally, because this tool is accessed from within the Cisco MDS 9000 Family switch, it is protected by the roles-based policy that limits access in each switch.

See the [“Capturing Frames Locally”](#) section on page 39-11.

Send documentation comments to mdsfeedback-doc@cisco.com.

Remote Capture Daemon

This daemon is the server end of the remote capture component. The Ethereal analyzer running on a host is the client end. They communicate with each other using the Remote Capture Protocol (RPCAP). RPCAP uses two endpoints, a TCP-based control connection and a TCP or UDP-based data connection based on TCP (default) or UDP. The control connection is used to remotely control the captures (start or stop the capture, or specify capture filters). Remote capture can only be performed to explicitly configured hosts. This technique prevents an unauthorized machine in the network from snooping on the control traffic in the network.

RPCAP supports two setup connection modes based on firewall restrictions.

- **Passive mode (default)**—The configured host initiates connection to the switch. Multiple hosts can be configured to be in passive mode and multiple hosts can be connected and receive remote captures at the same time.
- **Active mode**—The switch initiates the connection to a configured host—one host at a time.

Using capture filters, you can limit the amount of traffic that is actually sent to the client. Capture filters are specified at the client end—on Ethereal, not on the switch.

See the “[Sending Captures to Remote IP Addresses](#)” section on page 39-12.

GUI-Based Client

The Ethereal software runs on a host, such as a PC or workstation, and communicates with the remote capture daemon. This software is available in the public domain from <http://www.ethereal.com>. The Ethereal GUI front-end supports a rich interface such as a colorized display, graphical assists in defining filters, and specific frame searches. These features are documented on Ethereal’s website.

While remote capture through Ethereal supports capturing and decoding Fibre Channel frames from a Cisco MDS 9000 Family switch, the host running Ethereal does not require a Fibre Channel connection to the switch. The remote capture daemon running on the switch sends the captured frames over the out-of-band Ethernet management port. This capability allows you to capture and decode Fibre Channel frames from your desktop or laptop.

See the “[Displaying Captured Frames](#)” section on page 39-13.

Configuring the Cisco Fabric Analyzer

You can configure the Cisco Fabric Analyzer to perform one of two captures.

- **Local capture**—The command setting to enable a local capture cannot be saved to persistent storage or synchronized to standby. Launches the textual version on the fabric analyzer directly on the console screen. The capture can also be saved on the local file system.
- **Remote capture**—The command setting to enable a remote capture can be saved to persistent storage. It can be synchronized to the standby supervisor module and a stateless restart can be issued, if required.

To use the Cisco Fabric Analyzer feature, traffic should be flowing to or from the supervisor module.

Send documentation comments to mdsfeedback-doc@cisco.com.


Capturing Frames Locally

To capture frames locally, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| | Note The options within Step 2 may be performed in any order. | |
| Step 2 | switch(config)# fc analyzer local Capturing on eth2 switch(config)# | Begins capturing the frames locally (supervisor module). |
| | switch(config)# fc analyzer local brief Capturing on eth2 switch(config)# | Displays the protocol summary in a brief format. |
| | switch(config)# fc analyzer local display-filter SampleF Capturing on eth2 | Displays the filtered frames. |
| | switch(config)# fc analyzer local limit-frame-size 64 Capturing on eth2 switch(config)# | Limits the size of the frame capture to the first 64 bytes. The allowed range is 64 to 65536 bytes. |
| | switch(config)# fc analyzer local limit-captured-frames 10 Capturing on eth2 switch(config)# | Limits the number of frames captured to 10. The allowed range is 0 to 2147483647 frames and the default is 100 frames. Use 0 if you do not want to limit the number of captured frames. |
| | Note Press Ctrl-c to stop a capture. Otherwise, the capture stops automatically after capturing 100 frames. You can change this default using the fc analyzer local limit-captured-frames number command. | |
| Step 3 | switch(config)# fc analyzer local write volatile:sample Capturing on eth2 switch(config)# | Saves the captured frames to a specified file (sample) in the volatile: directory. Note Optionally, you can save the specified file to the slot0: directory. |
| | Note The final filename that is the capture file is called either SampleFile_00000_<dateandtime> or SampleFile_00001_<dateandtime>. For example, "SampleFile_00000_20021110223833" or "SampleFile_00001_20021110243833". The maximum size of a file that can be written to is 10 MB. | |

Send documentation comments to mdsfeedback-doc@cisco.com.

Sending Captures to Remote IP Addresses



Caution

You must use the eth2 interface to capture control traffic on a supervisor module.

To capture frames remotely, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fcanalyzer remote 10.21.0.3 switch(config)# | Configures the remote IP address (10.21.0.3) to which the captured frames are sent. |
| | switch(config)# fcanalyzer remote 10.21.0.3 active switch(config)# | Enables active mode (passive is the default) with the remote host. Ethereal is assumed to be running when the capture is performed. The switch tries to connect forever unless a capture stop instruction is sent from Ethereal. |
| | switch(config)# fcanalyzer remote 10.21.0.3 active 1 switch(config)# | Enables the active mode for a specified port. The valid port range is 1 to 65535. |

To capture remote traffic, use one of the following options:

- The capture interface can be specified in Ethereal as the remote device:

```
rpcap://<ipaddress or switch hostname>/eth2
```

For example:

```
rpcap://cp-16/eth2
rpcap://17.2.1.1/eth2
```
- The capture interface can be specified either in the capture dialog box or by using the -i option at the command line when invoking Ethereal.

```
ethereal -i rpcap://<ipaddress|hostname>[:<port>]/<interface>
```

For example:

```
ethereal -i rpcap://172.22.1.1/eth2
```

or

```
ethereal -i rpcap://customer-switch.customer.com/eth2
```



Note For example, in a Windows 2000 setup, click **Start** on your desktop and select **Run**. In the resulting Run window, type the required command line option in the Open field.

Clearing Configured fcanalyzer Information

Use the **clear fcanalyzer** command to clear the entire list of configured hosts. Note that the existing connections are not terminated.

Send documentation comments to mdsfeedback-doc@cisco.com.

Displaying Configured Hosts

Use the **show fcanalyzer** command to display the list of hosts configured for a remote capture. See [Example 39-3](#).

Example 39-3 Displays Configured Hosts

```
switch# show fcanalyzer
PassiveClient = 10.21.0.3
PassiveClient = 10.21.0.3
ActiveClient = 10.21.0.3, DEFAULT
```

**Note**

The DEFAULT in the ActiveClient line indicates that the default port is used.

Displaying Captured Frames

You can selectively view captured frames by using the display filters feature. For example, instead of viewing all the frames from a capture, you may only want to view Exchange Link Protocol (ELP) request frames. This feature only limits the captured view—it does not affect the captured or the saved frames. Procedures to specify, use, and save display filters are already documented in the [Ethereal website \(http://www.ethereal.com\)](http://www.ethereal.com). Some examples of how you can use this feature are as follows:

- To view all packets in a specified VSAN, use this expression:

```
mdshdr.vsan == 2
```

- To view all SW_ILS frames, use this expression:

```
fcswils
```

- To view class F frames, use this expression:

```
mdshdr.sof == SOFf
```

- To view all FSPF frames, use this expression:

```
swils.opcode == HLO || swils.opcode == LSU || swils.opcode == LSA
```

- To view all FLOGI frames, use this expression:

```
fcels.opcode == FLOGI
```

- To view all FLOGI frames in VSAN 1, use this expression:

```
fcels.opcode == FLOGI && mdshdr.vsan == 2
```

- To view all name server frames, use this expression:

```
dns
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Defining Display Filters

Display filters limit the frames that can be displayed, but not what is captured (similar to any view command). The filters to be displayed can be defined in multiple ways in the GUI application:

- Auto-definition
- Manual definition
- Assisted manual definition
- Only manual definition in local capture
- No assists

Regardless of the definition, each filter must be saved and identified with a name.



Note

This GUI-assisted feature is part of Ethereal and you can obtain more information from <http://www.ethereal.com>.

Displaying Filters Examples

Some examples of using display filters with the Fabric Analyzer local are provided in this section. The **brief** option is used in all examples to restrict the size of the output. See [Example 39-4](#).

Example 39-4 Displays Only Fabric Login Server Traffic on VSAN 1

```
switch(config)# fcanalyzer local brief display-filter
(mdshdr.vsan==0x01)&&((fc.d_id=="ff.ff.fe"\\|fc.s_id=="ff.ff.fe"))
Capturing on eth2
8.904145 00.00.00 -> ff.ff.fe FC ELS 1 0x28f8 0xffff 0x3 -> 0xf FLOGI
8.918164 ff.ff.fe -> 79.03.00 FC ELS 1 0x28f8 0x12c6 0xff -> 0x0 ACC (FLOGI)
```

You can trace all frames to and from a particular N port device. For example, to observe RSCNs from the Fabric Controller and registration and/or query requests to the name server. See [Example 39-5](#).



Note

The filter requires prior knowledge of the FC ID that is assigned to the N port. Issue the **show flogi database interface** command before running fcanalyzer to obtain the FC ID. In this example, the N port FC ID is 79.03.00.

Example 39-5 Displays All Traffic for a Particular N Port on VSAN 1

```
switch(config)# fcanalyzer local brief
display-filter(mdshdr.vsan==0x01)&&((fc.d_id=="79.03.00"\\|fc.s_id=="79.03.00"))
Capturing on eth2
8.699162 ff.ff.fe -> 79.03.00 FC ELS 1 0x35b8 0x148e 0xff -> 0x0 ACC (FLOGI)
8.699397 79.03.00 -> ff.ff.fc FC ELS 1 0x35d0 0xffff 0x3 -> 0xf PLOGI
8.699538 ff.ff.fc -> 79.03.00 FC ELS 1 0x35d0 0x148f 0xff -> 0x0 ACC (PLOGI)
8.699406 79.03.00 -> ff.ff.fd FC ELS 1 0x35e8 0xffff 0x3 -> 0xf SCR
8.700179 79.03.00 -> ff.ff.fc dNS 1 0x3600 0xffff 0x3 -> 0xf GNN_FT
8.702446 ff.ff.fd -> 79.03.00 FC ELS 1 0x35e8 0x1490 0xff -> 0x0 ACC (SCR)
8.704210 ff.ff.fc -> 79.03.00 dNS 1 0x3600 0x1491 0xff -> 0x0 ACC (GNN_FT)
8.704383 79.03.00 -> ff.ff.fc dNS 1 0x3618 0xffff 0x3 -> 0xf GPN_ID
8.707857 ff.ff.fc -> 79.03.00 dNS 1 0x3618 0x1496 0xff -> 0x0 ACC (GPN_ID)
```

Send documentation comments to mdsfeedback-doc@cisco.com.

The VSAN ID is specified in hex. See [Example 39-6](#).

Example 39-6 Displays All Traffic for a Specified VSAN

```
switch(config)# fcanalyzer local brief display-filter mdshdr.vsan==0x03e7
Capturing on eth2
12.762577 ff.ff.fd -> ff.ff.fd SW_ILS 999 0xb2c 0xffff 0x1 -> 0xf HLO
12.762639 ff.ff.fd -> ff.ff.fd FC 999 0xb2c 0xd32 0xff -> 0x0 Link Ctl, ACK1
13.509979 ff.ff.fd -> ff.ff.fd SW_ILS 999 0xd33 0xffff 0xff -> 0x0 HLO
13.510918 ff.ff.fd -> ff.ff.fd FC 999 0xd33 0xb2d 0x1 -> 0xf Link Ctl, ACK1
14.502391 ff.fc.64 -> ff.fc.70 SW_ILS 999 0xd34 0xffff 0xff -> 0x0 SW_RSCN
14.502545 ff.ff.fd -> 64.01.01 FC ELS 999 0xd35 0xffff 0xff -> 0x0 RSCN
14.502804 64.01.01 -> ff.ff.fd FC ELS 999 0xd35 0x215 0x0 -> 0xf ACC (RSCN)
14.503387 ff.fc.70 -> ff.fc.64 FC 999 0xd34 0xb2e 0x1 -> 0xf Link Ctl, ACK1
14.503976 ff.fc.70 -> ff.fc.64 SW_ILS 999 0xd34 0xb2e 0x1 -> 0xf SW_ACC (SW_RSCN)
14.504025 ff.fc.64 -> ff.fc.70 FC 999 0xd34 0xb2e 0xff -> 0x0 Link Ctl, ACK1
```

By excluding FSPF hellos and ACK1, you can focus on the frames of interest. See [Example 39-7](#).

Example 39-7 Displays All VSAN 1 Traffic Excluding FSPF Hellos and ACK1 Frames.

```
switch(config)# fcan lo bri dis
(mdshdr.vsan==0x01)&&not((swils.opcode==0x14)or(fc.r_ctl==0xc0))
Capturing on eth2
10.589934 ff.fc.79 -> ff.fc.7a FC-FCS 1 0x1b23 0xffff 0xff -> 0x0 GCAP
10.591253 ff.fc.7a -> ff.fc.79 FC-FCS 1 0x1b23 0x2f70 0x4 -> 0xf MSG_RJT (GCAP)
25.277981 ff.fc.79 -> ff.fc.7a SW_ILS 1 0x1b27 0xffff 0xff -> 0x0 SW_RSCN
25.278050 ff.fc.79 -> ff.fc.89 SW_ILS 1 0x1b28 0xffff 0xff -> 0x0 SW_RSCN
25.279232 ff.fc.89 -> ff.fc.79 SW_ILS 1 0x1b28 0xadd7 0x5 -> 0xf SW_ACC (SW_RSCN)
25.280023 ff.fc.7a -> ff.fc.79 Unzoned NS 1 0x3b2b 0xffff 0x5 -> 0xf GE_PT
25.280029 ff.fc.7a -> ff.fc.79 SW_ILS 1 0x1b27 0x2f71 0x4 -> 0xf SW_ACC (SW_RSCN)
25.282439 ff.fc.79 -> ff.fc.7a dNS 1 0x3b2b 0x1b29 0xff -> 0x0 RJT (GE_PT)
38.249966 00.00.00 -> ff.ff.fe FC ELS 1 0x36f0 0xffff 0x3 -> 0xf FLOGI
38.262622 ff.ff.fe -> 79.03.00 FC ELS 1 0x36f0 0x1b2b 0xff -> 0x0 ACC (FLOGI)
38.262844 79.03.00 -> ff.ff.fc FC ELS 1 0x3708 0xffff 0x3 -> 0xf PLOGI
38.262984 ff.ff.fc -> 79.03.00 FC ELS 1 0x3708 0x1b2c 0xff -> 0x0 ACC (PLOGI)
38.262851 79.03.00 -> ff.ff.fd FC ELS 1 0x3720 0xffff 0x3 -> 0xf SCR
38.263514 ff.fc.79 -> ff.fc.7a SW_ILS 1 0x1b2e 0xffff 0xff -> 0x0 SW_RSCN
38.263570 ff.fc.79 -> ff.fc.89 SW_ILS 1 0x1b2f 0xffff 0xff -> 0x0 SW_RSCN
38.263630 79.03.00 -> ff.ff.fc dNS 1 0x3738 0xffff 0x3 -> 0xf GNN_FT
38.263884 ff.ff.fd -> 79.03.00 FC ELS 1 0x3720 0x1b2d 0xff -> 0x0 ACC (SCR)
38.264066 ff.fc.89 -> ff.fc.79 SW_ILS 1 0x1b2f 0xaddf 0x5 -> 0xf SW_ACC (SW_RSCN)
38.264417 ff.fc.89 -> ff.fc.79 dNS 1 0xade0 0xffff 0x5 -> 0xf GE_ID
38.264585 ff.fc.79 -> ff.fc.89 dNS 1 0xade0 0x1b31 0xff -> 0x0 ACC (GE_ID)
38.265132 ff.ff.fc -> 79.03.00 dNS 1 0x3738 0x1b30 0xff -> 0x0 ACC (GNN_FT)
38.265210 ff.fc.7a -> ff.fc.79 Unzoned NS 1 0x3b2f 0xffff 0x5 -> 0xf GE_PT
38.265414 79.03.00 -> ff.ff.fc dNS 1 0x3750 0xffff 0x3 -> 0xf GPN_ID
38.265502 ff.fc.7a -> ff.fc.79 SW_ILS 1 0x1b2e 0x2f73 0x4 -> 0xf SW_ACC (SW_RSCN)
38.267196 ff.fc.79 -> ff.fc.7a dNS 1 0x3b2f 0x1b32 0xff -> 0x0 ACC (GE_PT)
```

Use this command to focus on TE port initialization. This example allows two VSANs on the TE port and the port VSAN is 666. Hence the ELP, ESC, and EPP (0x71) go out on VSAN 666. Once the EPP negotiation is complete, we see EFP, DIA, RDI, MR, FSPF, and other updates flow for each allowed VSAN. See [Example 39-8](#).

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 39-8 Displays SW_ILS Traffic Between Fabric Controllers for all VSANs and Exclude FSPF Hellos and ACK1 Frames.

```
switch(config)# fcan lo bri dis
fc.type==0x22&&((fc.d_id=="ff.fc.ef"\\|fc.s_id=="ff.fc.ef"))
Warning:Couldn't obtain netmask info (eth2:no IPv4 address assigned).
Capturing on eth2
9.472181 ff.fc.ef -> ff.fc.61 0x5e0a 0xffff SW_ILS ACA
9.472777 ff.fc.61 -> ff.fc.ef 0x5e0a 0x5e09 SW_ILS SW_ACC (ACA)
9.474551 ff.fc.ef -> ff.fc.61 0x5e0b 0xffff SW_ILS SFC
9.475706 ff.fc.61 -> ff.fc.ef 0x5e0b 0x5e0a SW_ILS SW_ACC (SFC)
9.476694 ff.fc.ef -> ff.fc.61 0x5e0c 0xffff SW_ILS UFC
9.483612 ff.fc.61 -> ff.fc.ef 0x5e0c 0x5e0b SW_ILS SW_ACC (UFC)
9.488187 ff.fc.ef -> ff.fc.61 0x5e0d 0xffff SW_ILS RCA
9.493703 ff.fc.61 -> ff.fc.ef 0x5e0d 0x5e0c SW_ILS SW_ACC (RCA)
```

This example focuses on zone server changes. Prior knowledge of the domain controller ID is required. The switch domain ID where the fcanalyzer is run is x79, the domain controller is FF.FC.79. See [Example 39-9](#).

Example 39-9 Display Switch Internal Link Services (SW_ILS) Traffic To and From Fabric Domain Controller ff.fc.79

```
switch(config)# fcan lo bri dis fc.type==0x22&&((fc.d_id=="ff.fc.79"\\|fc.s_id=="ff.fc.79"))
Capturing on eth2
64.053927 ff.fc.79 -> ff.fc.7a SW_ILS 0x1e15 0xffff 0xff -> 0x0 ACA
64.053995 ff.fc.79 -> ff.fc.89 SW_ILS 0x1e16 0xffff 0xff -> 0x0 ACA
64.054599 ff.fc.89 -> ff.fc.79 SW_ILS 0x1e16 0xb1e2 0x5 -> 0xf SW_ACC (ACA)
64.054747 ff.fc.7a -> ff.fc.79 SW_ILS 0x1e15 0x3037 0x4 -> 0xf SW_ACC (ACA)
64.057643 ff.fc.79 -> ff.fc.7a SW_ILS 0x1e17 0xffff 0xff -> 0x0 SFC
64.057696 ff.fc.79 -> ff.fc.89 SW_ILS 0x1e18 0xffff 0xff -> 0x0 SFC
64.058788 ff.fc.7a -> ff.fc.79 SW_ILS 0x1e17 0x3038 0x5 -> 0xf SW_ACC (SFC)
64.059288 ff.fc.89 -> ff.fc.79 SW_ILS 0x1e18 0xb1e3 0x5 -> 0xf SW_ACC (SFC)
64.062011 ff.fc.79 -> ff.fc.7a SW_ILS 0x1e19 0xffff 0xff -> 0x0 UFC
64.062060 ff.fc.79 -> ff.fc.89 SW_ILS 0x1e1a 0xffff 0xff -> 0x0 UFC
64.073513 ff.fc.7a -> ff.fc.79 SW_ILS 0x1e19 0x3039 0x5 -> 0xf SW_ACC (UFC)
64.765306 ff.fc.89 -> ff.fc.79 SW_ILS 0x1e1a 0xb1e4 0x5 -> 0xf SW_ACC (UFC)
64.765572 ff.fc.79 -> ff.fc.7a SW_ILS 0x1e1b 0xffff 0xff -> 0x0 RCA
64.765626 ff.fc.79 -> ff.fc.89 SW_ILS 0x1e1c 0xffff 0xff -> 0x0 RCA
64.766386 ff.fc.7a -> ff.fc.79 SW_ILS 0x1e1b 0x303a 0x4 -> 0xf SW_ACC (RCA)
64.766392 ff.fc.89 -> ff.fc.79 SW_ILS 0x1e1c 0xb1e5 0x5 -> 0xf SW_ACC (RCA)
```



Note

You can find the Fabric Domain Controller address in the Mgmt-Id field in the **show fcs ie vsan** command output.

```
switch# show fcs ie vsan 999
```

```
IE List for VSAN:999
```

| IE-WWN | IE-Type | Mgmt-Id | Mgmt-Addr |
|-------------------------|-----------------------|------------------|-------------|
| 23:e7:00:05:30:00:91:5f | Switch (Remote) | 0xffffc04 | 10.66.78.51 |
| 23:e7:00:05:30:00:9b:9f | Switch (Adjacent) | 0xffffc01 | 10.66.78.52 |
| 23:e7:00:0d:ec:00:93:81 | Switch (Local) | 0xffffc79 | 10.66.78.54 |

[Total 3 IEs in Fabric]

Send documentation comments to mdsfeedback-doc@cisco.com.

Capture Filters

You can limit what frames are captured by using the capture filters feature in a remote capture. This feature limits the frames that are captured and sent from the remote switch to the host. For example, you can capture only class F frames. Capture filters are useful in restricting the amount of bandwidth consumed by the remote capture.

Unlike display filters, capture filters restrict a capture to the specified frames. No other frames are visible until you specify a completely new capture.

The syntax for capture filter is different from the syntax for display filters. Capture filters use the Berkeley Packet Filter (BPF) library that is used in conjunction with the libpcap freeware. The list of all valid Fibre Channel capture filter fields are provided later in this section.

Procedures to configure capture filters are already documented in the Ethereal website (<http://www.ethereal.com>). Some examples of how you can use this feature as follows:

- To capture frames only on a specified VSAN, use this expression:

```
vsan = 1
```

- To capture only class F frames, use this expression:

```
class_f
```

- To capture only class Fibre Channel ELS frames, use this expression:

```
els
```

- To capture only name server frames, use this expression:

```
dns
```

- To capture only SCSI command frames, use this expression:

```
fcp_cmd
```



Note

This feature is part of libpcap and you can obtain more information from <http://www.tcpdump.org>.

Permitted Capture Filters

This section lists the permitted capture filters.

- o vsan
- o src_port_idx
- o dst_port_idx
- o sof
- o r_ctl
- o d_id
- o s_id
- o type
- o seq_id
- o seq_cnt
- o ox_id
- o rx_id
- o els
- o swils
- o fcp_cmd (FCP Command frames only)
- o fcp_data (FCP data frames only)
- o fcp_rsp (FCP response frames only)
- o class_f

Send documentation comments to mdsfeedback-doc@cisco.com.

```

o bad_fc
o els_cmd
o swils_cmd
o fcp_lun
o fcp_task_mgmt
o fcp_scsi_cmd
o fcp_status
o gs_type      (Generic Services type)
o gs_subtype   (Generic Services subtype)
o gs_cmd
o gs_reason
o gs_reason_expl
o dns      (name server)
o udns     (unzoned name server)
o fcs      (fabric configuration server)
o zs       (zone server)
o fc       (use as fc[x:y] where x is offset and y is length to compare)
o els      (use as els[x:y] similar to fc)
o swils    (use as swils[x:y] similar to fc)
o fcp      (use as fcp[x:y] similar to fc)
o fcct     (use as fcct[x:y] similar to fc)

```

Configuring World Wide Names

The world wide name (WWN) in the switch is equivalent to the Ethernet MAC address. As with the MAC address, you must uniquely associate the WWN to a single device. The principal switch selection and the allocation of domain IDs rely on the WWN. The WWN manager, a process-level manager residing on the switch’s supervisor module, assigns WWNs to each switch.

Cisco MDS 9000 Family switches support three network address authority (NAA) address formats (see [Table 39-1](#)).

Table 39-1 Standardized NAA WWN Formats

| NAA Address | NAA Type | WWN Format | |
|---------------------|----------------|--------------------------|--------------------|
| IEEE 48-bit address | Type 1 = 0001b | 000 0000 0000b | 48-bit MAC address |
| IEEE extended | Type 2 = 0010b | Locally assigned | 48-bit MAC address |
| IEEE registered | Type 5 = 0101b | IEEE company ID: 24 bits | VSID: 36 bits |



Caution

Changes to the world-wide names should be made by an administrator or individual who is completely familiar with switch operations.

Send documentation comments to mdsfeedback-doc@cisco.com.

Link Initialization WWN Usage

Exchange Link Protocol (ELP) and Exchange Fabric Protocol (EFP) use WWNs during link initialization. The usage details differ based on the Cisco SAN-OS software release:

- In Cisco SAN-OS Release 1.0 and 1.1, both ELPs and EFPs use the VSAN WWN during link initialization.
- In Cisco SAN-OS Releases 1.2 and 1.3, two different WWNs are used during the link initialization process:
 - ELPs use the switch WWN.
 - EFPs use the VSAN WWN.
- In Cisco SAN-OS Release 2.0(1b), both ELPs and EFPs use the VSAN WWN by default during link initialization. However, the ELP usage changes based on the peer switch's usage:
 - If the peer switch ELP uses the switch WWN, then the local switch also uses the switch WWN.
 - If the peer switch ELP uses the VSAN WWN, then the local switch also uses the VSAN WWN.

This link initialization change between Cisco SAN-OS releases is implicit and does not require any configuration.



Note

As of Cisco SAN-OS Release 2.0(2b), the ELP is enhanced to be compliant with FC-SW-3.

Configuring a Secondary MAC Address

To allocate secondary MAC addresses, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# wn secondary-mac 00:99:55:77:55:55 range 64 This command CANNOT be undone. Please enter the BASE MAC ADDRESS again: 00:99:55:77:55:55 Please enter the mac address RANGE again: 64 From now on WWN allocation would be based on new MACs. Are you sure? (yes/no) no You entered: no. Secondary MAC NOT programmed switch(config)# | Configures the secondary MAC address. This command cannot be undone. |

Displaying WWN Information

Use the **show wwn** commands to display the status of the WWN configuration. See Examples 39-10 to 39-12.

Example 39-10 Displays the Status of All WWNs

```
switch# show wwn status
      Type 1 WWNs: Configured:      64 Available:      48 (75%) Resvd.: 16
      Types 2 & 5 WWNs: Configured: 524288 Available: 450560 (85%) Resvd.: 73728
      NKAU & NKCR WWN Blks: Configured: 1760 Available: 1760 (100%)
      Alarm Status:      Type1:      NONE Types 2&5:      NONE
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 39-11 Displays Specified Block ID Information

```
switch# show wwn status block-id 51
WWNs in this block: 21:00:ac:16:5e:52:00:03 to 21:ff:ac:16:5e:52:00:03
Num. of WWNs:: Configured: 256 Allocated:    0 Available: 256
Block Allocation Status: FREE
```

Example 39-12 Displays the WWN for a Specific Switch

```
switch# show wwn switch
Switch WWN is 20:00:ac:16:5e:52:00:00
```

FC ID Allocation for HBAs

Fibre Channel standards require a unique FC ID to be allocated to an N port attached to a Fx port in any switch. To conserve the number of FC IDs used, Cisco MDS 9000 Family switches use a special allocation scheme.

Some HBAs do not discover targets that have FC IDs with the same domain and area. Prior to Cisco SAN-OS Release 2.0(1b), the Cisco SAN-OS software maintained a list of tested company IDs which do not exhibit this behavior. These HBAs were allocated with single FC IDs, and for others a full area was allocated.

The FC ID allocation scheme available in Release 1.3 and earlier, allocates a full area to these HBAs. This allocation isolates them to that area and are listed with their pWWN during a fabric login. The allocated FC IDs are cached persistently and are still available in Cisco SAN-OS Release 2.0(1b) (see the [“FC ID Allocation for HBAs” section on page 39-20](#)).

As of Cisco SAN-OS Release 2.0(1b), to allow further scalability for switches with numerous ports, the Cisco SAN-OS software is maintaining a list of HBAs exhibiting this behavior. Each HBA is identified by its company ID (also known as Organizational Unit Identifier, or OUI) used in the pWWN during a fabric log in. Hence a full area is allocated to the N ports with company IDs that are listed and for the others, a single FC ID is allocated. Irrespective of the kind (whole area or single) of FC ID allocated, the FC ID entries remain persistent.

Default Company ID list

All switches in the Cisco MDS 9000 Family, that ship with the SAN-OS Software Release 2.0(1b) or later, contain a default list of company IDs that require area allocation. Using the company ID reduces the number of configured persistent FC ID entries. You can configure or modify these entries using the CLI.



Caution

Persistent entries take precedence over company ID configuration. If the HBA fails to discover a target, verify that the HBA and the target are connected to the same switch and have the same area in their FC IDs, then perform the following procedure.

- Shut down the port connected to the HBA.
- Clear the persistent FC ID entry.
- Get the company ID from Port WWN.
- Add the company ID to the list that requires area.
- Bring up the port.

Send documentation comments to mdsfeedback-doc@cisco.com.

The list of company IDs have the following characteristics:

- A persistent FC ID configuration always takes precedence over the list of company IDs. Hence even if the company ID is configured to receive an area, the persistent FC ID configuration results in the allocation of a single FC ID.
- New company IDs added to subsequent releases are automatically added to existing company IDs.
- The list of company IDs is saved as part of the running and saved configuration.
- The list of company IDs is used only when the fcinterop FC ID allocation scheme is in **auto** mode. By default, the interop FC ID allocation is set to auto, unless changed.



Tip

We recommend that you set the fcinterop FC ID allocation scheme to **auto** and use the company ID list and persistent FC ID configuration to manipulate the FC ID device allocation.

Use the **fcinterop FCID allocation auto** command to change the FC ID allocation and the **show running-config** command to view the currently-allocated mode.

- When you issue a **write erase**, the list inherits the default list of company IDs shipped with a relevant release.

To allocate company IDs, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fcid-allocation area company-id 0x003223 | Adds a new company ID to the default list. |
| | switch(config)# no fcid-allocation area company-id 0x00E069 | Deletes a company ID from the default list. |
| | switch(config)# fcid-allocation area company-id 0x003223 | Adds a new company ID to the default list. |

Company ID Configuration Verification

You can view the configured company IDs by issuing the **show fcid-allocation area** command (see [Example 39-13](#)). Default entries are listed first and the user added entries are listed next. Entries are listed even if they were part of the default list and you later removed them.

Example 39-13 Displays the List of Default and Configured Company IDs

```
switch# show fcid-allocation area
FCID area allocation company id info:
  00:50:2E <----- Default entry
  00:50:8B
  00:60:B0
  00:A0:B8
  00:E0:69
  00:30:AE + <----- User-added entry
  00:32:23 +

  00:E0:8B * <----- Explicitly deleted entry (from the original default list)
Total company ids: 7
+ - Additional user configured company ids.
* - Explicitly deleted company ids from default list.
```

Send documentation comments to mdsfeedback-doc@cisco.com.

You can implicitly derive the default entries shipped with a specific release by combining the list of Company IDs displayed without any identification with the list of deleted entries.

You can also view or obtain the company IDs in a specific WWN by issuing the **show fcid-allocation company-id-from-wwn** command (see [Example 39-13](#)). Some WWN formats do not support company IDs. In these cases, you may need to configure the FC ID persistent entry.

Example 39-14 Displays the Company ID for the Specified WWN

```
switch# show fcid-allocation company-id-from-wwn 20:00:00:05:30:00:21:60
Extracted Company ID: 0x000530
```

Loop Monitoring Initiation

By default, the loop monitoring is disabled in all switches in the Cisco MDS 9000 Family. When a disk is removed from a loop port, the loop stays active based on the bypass circuit. Thus the disk removal is not known until you try to communicate with the disk. To detect such removals, the disks can be polled periodically (every 20 seconds).



Changes to the loop monitoring feature should be made by an administrator or individual who is completely familiar with switch operations.

Use the **fcinterop loop-monitor** command to enable loop polling for FL ports in a Cisco MDS 9000 Family switch.

To enable the loop monitoring feature, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fcinterop loop-monitor | Enables the loop monitoring feature. |
| | switch(config)# no fcinterop loop-monitor | Disables (default) the loop monitoring feature and reverts the switch to the factory defaults. |

Switch Interoperability

Interoperability enables the products of multiple vendors to come into contact with each other. Fibre Channel standards guide vendors towards common external Fibre Channel interfaces.

If all vendors followed the standards in the same manner, then interconnecting different products would become a trivial exercise. However, not all vendors follow the standards in the same way thus resulting in interoperability modes. This section briefly explains the basic concepts of these modes.

Each vendor has a regular mode and an equivalent interoperability mode, which specifically turns off advanced or proprietary features and provides the product with a more aimable standards compliant implementation.

[Table 39-2](#) lists the changes in switch behavior when you enable interoperability mode. These changes are specific to switches in the Cisco MDS 9000 Family while in interop mode.

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 39-2 *Changes in Switch Behavior When Interoperability Is Enabled*

| Switch Feature | Changes if Interoperability Is Enabled |
|--------------------------------------|---|
| Domain IDs | Some vendors cannot use the full range of 239 domains within a fabric. Domain IDs are restricted to the range 97-127. This is to accommodate McData's nominal restriction to this same range. They can either be set up statically (the Cisco MDS switch accept only one domain ID, if it does not get that domain ID it isolates itself from the fabric) or preferred. (If it does not get its requested domain ID, it accepts any assigned domain ID.) |
| Timers | All Fibre Channel timers must be the same on all switches as these values are exchanged by E ports when establishing an ISL. The timers are F_S_TOV, D_S_TOV, E_D_TOV, and R_A_TOV. |
| F_S_TOV | Verify that the Fabric Stability Time Out Value timers match exactly. |
| D_S_TOV | Verify that the Distributed Services Time Out Value timers match exactly. |
| E_D_TOV | Verify that the Error Detect Time Out Value timers match exactly. |
| R_A_TOV | Verify that the Resource Allocation Time Out Value timers match exactly. |
| Trunking | Trunking is not supported between two different vendor's switches. This feature may be disabled on a per port or per switch basis. |
| Default zone | The default zone behavior of permit (all nodes can see all other nodes) or deny (all nodes are isolated when not explicitly placed in a zone) may change. |
| Zoning attributes | Zones may be limited to the pWWN and other proprietary zoning methods (physical port number) may be eliminated. Note Brocade uses the cfgsave command to save fabric-wide zoning configuration. This command does not have any effect on Cisco MDS 9000 Family switches if they are part of the same fabric. You must explicitly save the configuration on each switch in the Cisco MDS 9000 Family. |
| Zone propagation | Some vendors do not pass the full zone configuration to other switches, only the active zone set gets passed. Verify that the active zone set or zone configuration has correctly propagated to the other switches in the fabric. |
| VSAN | Interop mode only affects the specified VSAN. |
| TE ports and PortChannels | TE ports and PortChannels cannot be used to connect Cisco MDS to non-Cisco MDS switches. Only E ports can be used to connect to non-Cisco MDS switches. TE ports and PortChannels can still be used to connect an Cisco MDS to other Cisco MDS switches even when in interop mode. |
| FSPF | The routing of frames within the fabric is not changed by the introduction of interop mode. The switch continues to use src-id, dst-id, and ox-id to load balance across multiple ISL links. |
| Domain reconfiguration disruptive | This is a switch-wide impacting event. Brocade and McData require the entire switch to be placed in offline mode and/or rebooted when changing domain IDs. |
| Domain reconfiguration nondisruptive | This event is limited to the affected VSAN. Only Cisco MDS 9000 Family switches have this capability—only the domain manager process for the affected VSAN is restarted and not the entire switch. |

Send documentation comments to mdsfeedback-doc@cisco.com.

Table 39-2 *Changes in Switch Behavior When Interoperability Is Enabled (continued)*

| Switch Feature | Changes if Interoperability Is Enabled |
|----------------|--|
| Name server | Verify that all vendors have the correct values in their respective name server database. |
| IVR | IVR-enabled VSANs can be configured in no interop (default) mode or in any of the interop modes. |

Configuring Interoperability

The interop mode in Cisco MDS 9000 Family switches can be enabled disruptively or nondisruptively.



Note

Brocade's `msplmgmtdeactivate` command must explicitly be run prior to connecting from a Brocade switch to either Cisco MDS 9000 Family switches or to McData switches. This command uses Brocade proprietary frames to exchange platform information, which Cisco MDS 9000 Family switches or McData switches do not understand. Rejecting these frames, causes the common E ports to become isolated.

To configure interoperability in any switch in the Cisco MDS 9000 Family, follow these steps:

- Step 1** Place the VSAN of the E ports (s) that connect to the OEM switch in interoperability mode.

```
switch# config t
switch(config)# vsan database
switch (config-vsan-db)# vsan 1 interop 1
```

- Step 2** Assign a domain ID in the range of 97 (0x61) through 127 (0x7F).



Note

This is an limitation imposed by the McData switches.

```
switch# config t
switch(config)# fcdomain domain 100 preferred vsan 1
```

In Cisco MDS 9000 switches, the default is to request an ID from the principal switch. If the **preferred** option is used, Cisco MDS 9000 switches request a specific ID, but still join the fabric if the principal switch assigns a different ID. If the **static** option is used, the Cisco MDS 9000 switches do not join the fabric unless the principal switch agrees, and assigns the requested ID.



Note

When changing the domain ID, the FC IDs assigned to N ports also change.

- Step 3** Change the Fibre Channel timers (if they have been changed from the system defaults).



Note

The Cisco MDS 9000, Brocade, and McData FC Error Detect (ED_TOV) and Resource Allocation (RA_TOV) timers default to the same values. They can be changed if needed. The RA_TOV default is 10 seconds, and the ED_TOV default is 2 seconds. Per the FC-SW2 standard, these values must be the same on each switch within the fabric.

Send documentation comments to mdsfeedback-doc@cisco.com.

```
switch# config t
switch(config)# fctimer e_d_tov ?
<1000-100000> E_D_TOV in milliseconds(1000-100000)
switch(config)# fctimer r_a_tov ?
<5000-100000> R_A_TOV in milliseconds(5000-100000)
```

Step 4 When making changes to the domain, you may or may not need to restart the Cisco MDS domain manager function for the altered VSAN.

- Force a fabric reconfiguration with the **disruptive** option.

```
switch(config)# fcdomain restart disruptive vsan 1
```

or

- Do not force a fabric reconfiguration.

```
switch(config)# fcdomain restart vsan 1
```

Verifying Interoperating Status

This section highlights the commands used to verify if the fabric is up and running in interoperability mode.

To verify the resulting status of issuing the interoperability command in any switch in the Cisco MDS 9000 Family, follow these steps:

Step 1 Use the **show version** command to verify the version.

```
switch# show version
Cisco Storage Area Networking Operating System (SAN-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2003, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
Cisco Systems, Inc. and/or other third parties and are used and
distributed under license. Some parts of this software are covered
under the GNU Public License. A copy of the license is available
at http://www.gnu.org/licenses/gpl.html.

Software
  BIOS:          version 1.0.8
  loader:        version 1.1(2)
  kickstart:     version 2.0(1) [build 2.0(0.6)] [gdb]
  system:        version 2.0(1) [build 2.0(0.6)] [gdb]

  BIOS compile time:      08/07/03
  kickstart image file is: bootflash:///m9500-sflek9-kickstart-mzg.2.0.0.6.bin
  kickstart compile time: 10/25/2010 12:00:00
  system image file is:   bootflash:///m9500-sflek9-mzg.2.0.0.6.bin
  system compile time:    10/25/2020 12:00:00

Hardware
  RAM 1024584 kB

  bootflash: 1000944 blocks (block size 512b)
  slot0:      0 blocks (block size 512b)

172.22.92.181 uptime is 0 days 2 hours 18 minute(s) 1 second(s)
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
Last reset at 970069 usecs after Tue Sep 16 22:31:25 1980
Reason: Reset Requested by CLI command reload
System version: 2.0(0.6)
Service:
```

Step 2 Use the **show interface brief** command to verify if the interface states are as required by your configuration.

```
switch# show int brief
Interface  Vsan    Admin  Admin  Status      Oper  Oper  Port-channel
          Mode    Trunk
          Mode
-----
fc2/1      1       auto   on     up           E     2     --
fc2/2      1       auto   on     up           E     2     --
fc2/3      1       auto   on     fcotAbsent  --    --    --
fc2/4      1       auto   on     down        --    --    --
fc2/5      1       auto   on     down        --    --    --
fc2/6      1       auto   on     down        --    --    --
fc2/7      1       auto   on     up           E     1     --
fc2/8      1       auto   on     fcotAbsent  --    --    --
fc2/9      1       auto   on     down        --    --    --
fc2/10     1       auto   on     down        --    --    --
```

Step 3 Use the **show run** command to verify if you are running the desired configuration.

```
switch# show run
Building Configuration...

interface fc2/1
no shutdown

interface fc2/2
no shutdown

interface fc2/3
interface fc2/4
interface fc2/5
interface fc2/6
interface fc2/7
no shutdown

interface fc2/8
interface fc2/9
interface fc2/10

<snip>

interface fc2/32

interface mgmt0
ip address 6.1.1.96 255.255.255.0
switchport encap default
no shutdown

vsan database
vsan 1 interop

boot system bootflash:/m9500-system-253e.bin sup-1
boot kickstart bootflash:/m9500-kickstart-253e.bin sup-1
boot system bootflash:/m9500-system-253e.bin sup-2
boot kickstart bootflash:/m9500-kickstart-253e.bin sup-2
callhome
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```

fcdomain domain 100 preferred vsan 1

ip route 6.1.1.0 255.255.255.0 6.1.1.1
ip routing
line console
  databits 5
  speed 110
logging linecard
ssh key rsa 512 force
ssh server enable
switchname MDS9509
username admin password 5 $1$Li8/fBYX$SNc72.xt4nTXpSnR9OUFB/ role network-admin

```

Step 4 Use the **show vsan** command to verify if the interoperability mode is active.

```

switch# show vsan 1
vsan 1 information
  name:VSAN0001 stalactites
  interoperability mode:yes <----- verify mode
  loadbalancing:src-id/dst-id/oxid
  operational state:up

```

Step 5 Use the **show fcdomain vsan** command to verify the domain ID.

```

switch# show fcdomain vsan 1
The local switch is a Subordinated Switch.

Local switch run time information:
  State: Stable
  Local switch WWN:      20:01:00:05:30:00:51:1f
  Running fabric name:  10:00:00:60:69:22:32:91
  Running priority: 128
  Current domain ID: 0x64(100) <-----verify domain id

Local switch configuration information:
  State: Enabled
  Auto-reconfiguration: Disabled
  Contiguous-allocation: Disabled
  Configured fabric name: 41:6e:64:69:61:6d:6f:21
  Configured priority: 128
  Configured domain ID: 0x64(100) (preferred)

Principal switch run time information:
  Running priority: 2

```

| Interface | Role | RCF-reject |
|-----------|------------|------------|
| fc2/1 | Downstream | Disabled |
| fc2/2 | Downstream | Disabled |
| fc2/7 | Upstream | Disabled |

Send documentation comments to mdsfeedback-doc@cisco.com.

Step 6 Use the **show fcdomain domain-list vsan** command to verify the local principal switch status.

```
switch# show fcdomain domain-list vsan 1

Number of domains: 5
Domain ID          WWN
-----
0x61(97)           10:00:00:60:69:50:0c:fe
0x62(98)           20:01:00:05:30:00:47:9f
0x63(99)           10:00:00:60:69:c0:0c:1d
0x64(100)          20:01:00:05:30:00:51:1f [Local]
0x65(101)          10:00:00:60:69:22:32:91 [Principal]
-----
```

Step 7 Use the **show fspf internal route vsan** command to verify the next hop and destination for the switch.

```
switch# show fspf internal route vsan 1

FSPF Unicast Routes
-----
VSAN Number  Dest Domain  Route Cost  Next hops
-----
              1      0x61(97)      500        fc2/2
              1      0x62(98)     1000        fc2/1
                      fc2/2
              1      0x63(99)      500        fc2/1
              1      0x65(101)     1000        fc2/7
```

Step 8 Use the **show fcns data vsan** command to verify the name server information.

```
switch# show fcns data vsan 1
VSAN 1:
-----
FCID          TYPE  PWWN                                (VENDOR) FC4-TYPE:FEATURE
-----
0x610400      N     10:00:00:00:c9:24:3d:90 (Emulex)   scsi-fcp
0x6105dc      NL    21:00:00:20:37:28:31:6d (Seagate)  scsi-fcp
0x6105e0      NL    21:00:00:20:37:28:24:7b (Seagate)  scsi-fcp
0x6105e1      NL    21:00:00:20:37:28:22:ea (Seagate)  scsi-fcp
0x6105e2      NL    21:00:00:20:37:28:2e:65 (Seagate)  scsi-fcp
0x6105e4      NL    21:00:00:20:37:28:26:0d (Seagate)  scsi-fcp
0x630400      N     10:00:00:00:c9:24:3f:75 (Emulex)   scsi-fcp
0x630500      N     50:06:01:60:88:02:90:cb (Seagate)  scsi-fcp
0x6514e2      NL    21:00:00:20:37:a7:ca:b7 (Seagate)  scsi-fcp
0x6514e4      NL    21:00:00:20:37:a7:c7:e0 (Seagate)  scsi-fcp
0x6514e8      NL    21:00:00:20:37:a7:c7:df (Seagate)  scsi-fcp
0x651500      N     10:00:00:e0:69:f0:43:9f (JNI)

Total number of entries = 12
```



Note

The Cisco MDS name server shows both local and remote entries, and does not time out the entries.

Send documentation comments to mdsfeedback-doc@cisco.com.

The show tech-support Command

The **show tech-support** command is useful when collecting a large amount of information about your switch for troubleshooting purposes. The output of this command can be provided to technical support representatives when reporting a problem.

The **show tech-support** command displays the output of several **show** commands at once. The output from this command varies depending on your configuration. Use the **show tech-support** command in EXEC mode to display general information about the switch when reporting a problem.

You can choose to have detailed information for each command or even specify the output for a particular interface, module or VSAN. Each command output is separated by line and the command precedes the output.



Note

Explicitly set the **terminal length** command to 0 (zero) to disable auto-scrolling and enable manual scrolling. Use the **show terminal** command to view the configured terminal size. After obtaining the output of this command, remember to reset your terminal length as required (see the [“Setting the Terminal Length”](#) section on page 2-18).



Tip

You can save the output of this command to a file by appending `> filename` to the **show tech-support** command (see the [“Saving Command Output to a File”](#) section on page 2-26). If you save this file, verify you have sufficient space to do so—each of these files may take about 1.8 MB. However, you can zip this file using the **gzip filename** command (see the [“Compressing and Uncompressing Files”](#) section on page 2-26). Copy the zipped file to the required location using the **copy** command and unzip the file using the **gunzip** command (see the [“Copying Files”](#) section on page 2-25).

The default output of the **show tech-support** command includes the output of the following commands:

- **show version**
- **show environment**
- **show module**
- **show hardware**
- **show running-config**
- **show interface**
- **show accounting log**
- **show process**
- **show process log**
- **show processes log details**
- **show flash**

Each command is discussed in both the *Cisco MDS 9000 Family Configuration Guide* and the *Cisco MDS 9000 Family Command Reference*. Refer to the *Cisco MDS 9000 Family Troubleshooting Guide* to obtain debug processes, procedures, and examples.

Send documentation comments to mdsfeedback-doc@cisco.com.

The show tech-support brief Command

Use the **show tech-support brief** command to obtain a quick, condensed review of your switch configurations. This command provides a summary of the current running state of the switch.

The **show tech-support brief** command is useful when collecting information about your switch for troubleshooting purposes. The output of this command can be provided to technical support representatives when reporting a problem.



Tip

You can save the output of this command to a file by appending `> filename` to the **show tech-support brief** command (see the “[Saving Command Output to a File](#)” section on page 2-26).

Example 39-15 Displays the Condensed View of Switch Configurations

```
vegas01# show tech-support brief
Switch Name           : vegas01
Switch Type           : DS-X9216-K9-SUP
Kickstart Image       : 1.3(2) bootflash:///m9200-ek9-kickstart-mz.1.3.1.10.bin
System Image          : 1.3(2) bootflash:///m9200-ek9-mz.1.3.1.10.bin
IP Address/Mask       : 10.76.100.164/24
Switch WWN            : 20:00:00:05:30:00:84:9e
No of VSANs           : 9
Configured VSANs      : 1-6,4091-4093

VSAN    1:    name:VSAN0001, state:active, interop mode:default
           domain id:0x6d(109), WWN:20:01:00:05:30:00:84:9f [Principal]
           active-zone:VR, default-zone:deny

VSAN    2:    name:VSAN0002, state:active, interop mode:default
           domain id:0x7d(125), WWN:20:02:00:05:30:00:84:9f [Principal]
           active-zone:<NONE>, default-zone:deny

VSAN    3:    name:VSAN0003, state:active, interop mode:default
           domain id:0xbe(190), WWN:20:03:00:05:30:00:84:9f [Principal]
           active-zone:<NONE>, default-zone:deny

VSAN    4:    name:VSAN0004, state:active, interop mode:default
           domain id:0x5a(90), WWN:20:04:00:05:30:00:84:9f [Principal]
           active-zone:<NONE>, default-zone:deny

VSAN    5:    name:VSAN0005, state:active, interop mode:default
           domain id:0x13(19), WWN:20:05:00:05:30:00:84:9f [Principal]
           active-zone:<NONE>, default-zone:deny

VSAN    6:    name:VSAN0006, state:active, interop mode:default
           domain id:0x1f(31), WWN:20:06:00:05:30:00:84:9f [Principal]
           active-zone:<NONE>, default-zone:deny

VSAN 4091:    name:VSAN4091, state:active, interop mode:default
           domain id:0x08(8), WWN:2f:fb:00:05:30:00:84:9f [Principal]
           active-zone:<NONE>, default-zone:deny

VSAN 4092:    name:VSAN4092, state:active, interop mode:default
           domain id:0x78(120), WWN:2f:fc:00:05:30:00:84:9f [Principal]
           active-zone:<NONE>, default-zone:deny

VSAN 4093:    name:VSAN4093, state:active, interop mode:default
           domain id:0x77(119), WWN:2f:fd:00:05:30:00:84:9f [Principal]
           active-zone:<NONE>, default-zone:deny
```

Send documentation comments to mdsfeedback-doc@cisco.com.

| Interface | Vsan | Admin Mode | Admin Trunk Mode | Status | FCOT | Oper Mode | Oper Speed (Gbps) | Port Channel |
|-----------|------|------------|------------------|--------------|------|-----------|-------------------|--------------|
| fc1/1 | 1 | auto | on | fcotAbsent | -- | -- | | -- |
| fc1/2 | 1 | auto | on | fcotAbsent | -- | -- | | -- |
| fc1/3 | 1 | auto | on | fcotAbsent | -- | -- | | -- |
| fc1/4 | 1 | auto | on | fcotAbsent | -- | -- | | -- |
| fc1/5 | 1 | auto | on | notConnected | sw1 | -- | | -- |
| fc1/6 | 1 | auto | on | fcotAbsent | -- | -- | | -- |
| fc1/7 | 1 | auto | on | fcotAbsent | -- | -- | | -- |
| fc1/8 | 1 | auto | on | fcotAbsent | -- | -- | | -- |
| fc1/9 | 1 | auto | on | fcotAbsent | -- | -- | | -- |
| fc1/10 | 1 | auto | on | fcotAbsent | -- | -- | | -- |
| fc1/11 | 1 | auto | on | fcotAbsent | -- | -- | | -- |
| fc1/12 | 1 | auto | on | fcotAbsent | -- | -- | | -- |
| fc1/13 | 1 | auto | on | fcotAbsent | -- | -- | | -- |
| fc1/14 | 1 | auto | on | fcotAbsent | -- | -- | | -- |
| fc1/15 | 1 | auto | on | fcotAbsent | -- | -- | | -- |
| fc1/16 | 1 | auto | on | fcotAbsent | -- | -- | | -- |

| Interface | Status | Speed (Gbps) |
|-----------|--------|--------------|
| sup-fc0 | up | 1 |

| Interface | Status | IP Address | Speed | MTU |
|-----------|--------|------------------|----------|------|
| mgmt0 | up | 10.76.100.164/24 | 100 Mbps | 1500 |

Default Settings

Table 39-3 lists the default settings for the features included in this chapter.

Table 39-3 Default Settings for Advanced Features

| Parameters | Default |
|--|----------------------|
| D_S_TOV | 5,000 milliseconds. |
| E_D_TOV | 2,000 milliseconds. |
| R_A_TOV | 10,000 milliseconds. |
| Timeout period to invoke fctrace | 5 seconds. |
| Number of frame sent by the fcping feature | 5 frames. |
| Remote capture connection protocol | TCP. |
| Remote capture connection mode | Passive. |
| Local capture frame limit s | 10 frames. |
| FC ID allocation mode | Auto mode. |
| Loop monitoring | Disabled. |

Send documentation comments to mdsfeedback-doc@cisco.com.



Configuring Fabric Configuration Servers

This chapter describes the Fabric Configuration Server (FCS) feature provided in the Cisco MDS 9000 Family of directors and switches. It includes the following sections:

- [About FCS, page 40-1](#)
- [FCS Name Specification, page 40-2](#)
- [Displaying FCS Information, page 40-3](#)
- [Default Settings, page 40-6](#)

About FCS

The Fabric Configuration Server (FCS) provides discovery of topology attributes and maintains a repository of configuration information of fabric elements. A management application is usually connected to the FCS on the switch through an N port. The FCS views the entire fabric based on the following objects:

- Interconnect element (IE) object—Each switch in the fabric corresponds to an IE object. One or more IE objects form a fabric.
- Port object—Each physical port in an IE corresponds to a port object. This includes the switch ports (xE, Fx, and TL ports) and its attached Nx ports.
- Platform object—A set of nodes may be defined as a platform object to make it a single manageable entity. These nodes are end-devices (host systems, storage subsystems) attached to the fabric. Platform objects reside at the edge switches of the fabric.

Each object has its own set of attributes and values. A null value may also be defined for some attributes.


In the Cisco MDS 9000 Family switch environment, multiple VSANs constitute a fabric, where one instance of the FCS is present per VSAN.

If you have attached a management application to a switch, all the frames directed towards the FCS in the switch are part of the port VSAN in the switch port (Fx port). Hence your view of the management application is limited only to this VSAN. However, information about other VSANs that this switch is part of can be obtained either through the SNMP or CLI.

In [Figure 40-1](#) Management Application 1 (M1) is connected through an F port with port VSAN ID 1 and Management Application 2 (M2) is connected through an F port with port VSAN ID 2. M1 can query the FCS information of switches S1 and S3, and M2 can query switches S3 and S4. Switch S2 information is not known to both of them. FCS operations can be done only on those switches that are visible in the VSAN. Note that M2 can send FCS requests only for VSAN 2 even though S3 is also a part of VSAN 1.

Send documentation comments to mdsfeedback-doc@cisco.com.

Figure 40-1 FCSs in a VSAN Environment



Significance of FCS

This section lists the significance of FCSs.

- FCSs support network management including the following:
 - N port management application can query and obtain information about fabric elements.
 - A SNMP Manager can use the FCS management information base (MIB) to start discovery and obtain information about the fabric topology.
- FCSs support TE and TL ports in addition to the standard F and E ports.
- FCS can maintain a group of modes with a logical name and management address when a platform registers with it. FCSs maintain a backup of all registrations in secondary storage and update it with every change. When a restart or switchover happens, FCSs retrieve the secondary storage information and rebuild its database.
- SNMP manager can query FCSs for all IEs, ports, and platforms in the fabric.

FCS Name Specification

You can specify if the unique name verification is for the entire fabric (globally) or only for locally (default) registered platforms.



Note

Set this command globally only if all switches in the fabric belong to the Cisco MDS 9000 Family.

Send documentation comments to mdsfeedback-doc@cisco.com.

To enable global checking of the platform name, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fcs plat-check-global vsan 1 switch(config)# no fcs plat-check-global vsan 1 | Enables global checking of platform name. Disables (default) global checking of platform name. |

To register platform attributes, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fcs register switch(config-fcs-register)# | Enters the FCS registration submodule. |
| Step 3 | switch(config-fcs-register)# platform name SamplePlatform vsan 1 switch(config-fcs-register-attrib)# switch(config-fcs-register)# no platform name SamplePlatform vsan 1 switch(config-fcs-register)# | Enters the FCS registration attributes submodule. Deletes a registered platform. |
| Step 4 | switch(config-fcs-register-attrib)# mgmt-addr 1.1.1.1 switch(config-fcs-register)# no mgmt-addr 1.1.1.1 | Configures the platform management address. Deletes all management addresses on the platform. |
| Step 5 | switch(config-fcs-register-attrib)# nwn 11:22:33:44:55:66:77:88 switch(config-fcs-register)# no nwn 11:22:33:44:55:66:77:88 | Configures the platform node name. Deletes the platform node name. |
| Step 6 | switch(config-fcs-register-attrib)# type 5 switch(config-fcs-register)# no type 5 | Configures the fc-gs-3 defined platform type. Deletes the configured type and reverts the switch to its factory default of unknown type. |
| Step 7 | switch(config-fcs-register-attrib)# exit | Exits the FCS registration attributes submodule. |
| Step 8 | switch(config-fcs-register)# exit switch(config)# | Exits the FCS registration submodule. |

Displaying FCS Information

Use the **show fcs** commands to display the status of the WWN configuration (see Example 40-1 to 40-9).

Example 40-1 Displays FCS Local Database Information

```
switch# show fcs database
FCS Local Database in VSAN: 1
-----
Switch WWN                : 20:01:00:05:30:00:16:df
Switch Domain Id          : 0x7f(127)
Switch Mgmt-Addresses     : snmp://172.22.92.58/eth-ip
                          : http://172.22.92.58/eth-ip
```

Send documentation comments to mdsfeedback-doc@cisco.com.

```
Fabric-Name           : 20:01:00:05:30:00:16:df
Switch Logical-Name   : 172.22.92.58
Switch Information List : [Cisco Systems*DS-C9509*0*20:00:00:05:30:00
Switch Ports:
-----
Interface  pWWN                                Type      Attached-pWWNs
-----
fc2/1      20:41:00:05:30:00:16:de  TE        20:01:00:05:30:00:20:de
fc2/2      20:42:00:05:30:00:16:de  Unknown   None
fc2/17     20:51:00:05:30:00:16:de  TE        20:0a:00:05:30:00:20:de

FCS Local Database in VSAN: 5
-----
Switch WWN           : 20:05:00:05:30:00:12:5f
Switch Domain Id     : 0xef(239)
Switch Mgmt-Addresses : http://172.22.90.171/eth-ip
                     : snmp://172.22.90.171/eth-ip
                     : http://10.10.15.10/vsan-ip
                     : snmp://10.10.15.10/vsan-ip
Fabric-Name          : 20:05:00:05:30:00:12:5f
Switch Logical-Name   : 172.22.90.171
Switch Information List : [Cisco Systems*DS-C9509**20:00:00:05:30:00:12:5e]
Switch Ports:
-----
Interface  pWWN                                Type      Attached-pWWNs
-----
fc3/1      20:81:00:05:30:00:12:5e  TE        22:01:00:05:30:00:12:9e
fc3/2      20:82:00:05:30:00:12:5e  TE        22:02:00:05:30:00:12:9e
fc3/3      20:83:00:05:30:00:12:5e  TE        22:03:00:05:30:00:12:9e
```

Example 40-2 Displays a List of All IEs for a Specific VSAN

```
switch# show fcs ie vsan 1
IE List for VSAN: 1
-----
IE-WWN           IE-Type           Mgmt-Id
-----
20:01:00:05:30:00:16:df  Switch (Local)      0xffffc7f
20:01:00:05:30:00:20:df  Switch (Adjacent)   0xffffc64
[Total 2 IEs in Fabric]
```

Example 40-3 Displays Interconnect Element Object Information for a Specific nWWN

```
switch# show fcs ie nwwn 20:01:00:05:30:00:16:df vsan 1
IE Attributes
-----
Domain-Id = 0x7f(127)
Management-Id = 0xffffc7f
Fabric-Name = 20:01:00:05:30:00:16:df
Logical-Name = 172.22.92.58
Management Address List =
    snmp://172.22.92.58/eth-ip
    http://172.22.92.58/eth-ip
Information List:
    Vendor-Name = Cisco Systems
    Model Name/Number = DS-C9509
    Release-Code = 0
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 40-4 Displays Information for a Specific Platform

```
switch# show fcs platform name SamplePlatform vsan 1
Platform Attributes
-----
Platform Node Names:
      11:22:33:44:55:66:77:88
Platform Type = Gateway
Platform Management Addresses:
      1.1.1.1
```

Example 40-5 Displays a List of Platforms for a Specified VSAN

```
switch# show fcs platform vsan 1
Platform List for VSAN: 1
Platform-Names
-----
SamplePlatform
[Total 1 Platforms in Fabric]
```

Example 40-6 Displays a List of Switch Ports in a Specified VSAN

```
switch# show fcs port vsan 24
Port List in VSAN: 24
      -- IE WWN: 20:18:00:05:30:00:16:df --
-----
Port-WWN                Type      Module-Type      Tx-Type
-----
20:41:00:05:30:00:16:de  TE_Port  SFP with Serial Id  Shortwave Laser
20:51:00:05:30:00:16:de  TE_Port  SFP with Serial Id  Shortwave Laser
[Total 2 switch-ports in IE]
      -- IE WWN: 20:18:00:05:30:00:20:df --
-----
Port-WWN                Type      Module-Type      Tx-Type
-----
20:01:00:05:30:00:20:de  TE_Port  SFP with Serial Id  Shortwave Laser
20:0a:00:05:30:00:20:de  TE_Port  SFP with Serial Id  Shortwave Laser
[Total 2 switch-ports in IE]
```

Example 40-7 Displays Port Information for a Specified pWWN

```
switch# show fcs port pwwn 20:51:00:05:30:00:16:de vsan 24
Port Attributes
-----
Port Type = TE_Port
Port Number = 0x1090000
Attached-Port-WWNs:
      20:0a:00:05:30:00:20:de
Port State = Online
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 40-8 Displays FCS Statistics

```
switch# show fcs statistics
FCS Statistics for VSAN: 1
-----
FCS Rx Get Reqs      :2
FCS Tx Get Reqs      :7
FCS Rx Reg Reqs      :0
FCS Tx Reg Reqs      :0
FCS Rx Dereg Reqs    :0
FCS Tx Dereg Reqs    :0
FCS Rx RSCNs         :0
...
FCS Statistics for VSAN: 30
-----
FCS Rx Get Reqs      :2
FCS Tx Get Reqs      :2
FCS Rx Reg Reqs      :0
FCS Tx Reg Reqs      :0
FCS Rx Dereg Reqs    :0
FCS Tx Dereg Reqs    :0
FCS Rx RSCNs         :0
FCS Tx RSCNs         :0
...
```

Example 40-9 Displays Platform Settings for Each VSAN

```
switch# show fcs vsan
-----
VSAN      Plat Check fabric-wide
-----
0001      Yes
0010      No
0020      No
0021      No
0030      No
```

Default Settings

Table 40-1 lists the default FCS settings.

Table 40-1 Default FCS Settings

| Parameters | Default |
|--------------------------------------|----------|
| Global checking of the platform name | Disabled |
| Platform node type | Unknown |



Monitoring System Processes and Logs

This chapter provides details on monitoring the health of the switch. It includes the following sections:

- [Displaying System Processes, page 41-1](#)
- [Displaying System Status, page 41-4](#)
- [Core and Log Files, page 41-6](#)
- [Kernel Core Dumps, page 41-8](#)
- [Online System Health Management, page 41-10](#)
- [Default Settings, page 41-18](#)

Displaying System Processes

Use the **show processes** command to obtain general information about all processes (see [Example 41-1](#) to [Example 41-6](#)).

Example 41-1 Displays System Processes

```
switch# show processes
PID      State  PC      Start_cnt  TTY  Process
-----
  868     S    2ae4f33e      1     -   snmpd
  869     S    2acee33e      1     -   rscn
  870     S    2ac36c24      1     -   qos
  871     S    2ac44c24      1     -   port-channel
  872     S    2ac7a33e      1     -   ntp
    -    ER      -      1     -   mdog
    -    NR      -      0     -   vbuilder
```

Where:

- PID = process ID.
- State = process state.
 - D = uninterruptible sleep (usually I/O).
 - R = runnable (on run queue).
 - S = sleeping.
 - T = traced or stopped.
 - Z = defunct (“zombie”) process.

Send documentation comments to mdsfeedback-doc@cisco.com.

- NR = not running.
- ER = should be running but currently not-running.
- PC = current program counter in hex format.
- Start_cnt = number of times a process has been started (or restarted).
- TTY = terminal that controls the process. A hyphen usually means a daemon not running on any particular TTY.
- Process = name of the process.

Example 41-2 Displays CPU Utilization Information

```
switch# show processes cpu
PID      Runtime(ms)   Invoked    uSecs   1Sec   Process
-----
 842      3807        137001      27     0.0   sysmgr
1112      1220         67974      17     0.0   syslogd
1269       220         13568      16     0.0   fcfwd
1276      2901         15419     188     0.0   zone
1277       738         21010      35     0.0   xbar_client
1278      1159         6789       170     0.0   wwn
1279       515         67617       7     0.0   vsan
```

Where:

- Runtime (ms) = CPU time the process has used, expressed in milliseconds.
- Invoked = number of times the process has been invoked.
- uSecs = microseconds of CPU time on average for each process invocation.
- 1Sec = CPU utilization in percentage for the last one second.

Example 41-3 Displays Process Log Information

```
switch# show processes log
Process      PID      Normal-exit  Stack-trace  Core      Log-create-time
-----
fspf         1339           N             Y           N   Jan  5 04:25
lcm          1559           N             Y           N   Jan  2 04:49
rib          1741           N             Y           N   Jan  1 06:05
```

Where:

- Normal-exit = whether or not the process exited normally.
- Stack-trace = whether or not there is a stack trace in the log.
- Core = whether or not there exists a core file.
- Log-create-time = when the log file got generated.

Example 41-4 Displays Detail Log Information About a Process

```
switch# show processes log pid 1339
Service: fspf
Description: FSPF Routing Protocol Application

Started at Sat Jan  5 03:23:44 1980 (545631 us)
Stopped at Sat Jan  5 04:25:57 1980 (819598 us)
Uptime: 1 hours 2 minutes 2 seconds
```


Send documentation comments to mdsfeedback-doc@cisco.com.

```
Start type: SRV_OPTION_RESTART_STATELESS (23)
Death reason: SYSMGR_DEATH_REASON_FAILURE_SIGNAL (2)
Exit code: signal 9 (no core)
CWD: /var/sysmgr/work
```

Virtual Memory:

```
CODE      08048000 - 0809A100
DATA      0809B100 - 0809B65C
BRK       0809D988 - 080CD000
STACK     7FFFFD20
TOTAL     23764 KB
```

Register Set:

```
EBX 00000005      ECX 7FFFF8CC      EDX 00000000
ESI 00000000      EDI 7FFFF6CC      EBP 7FFFF95C
EAX FFFFFFFD     XDS 8010002B      XES 0000002B
EAX 0000000E (orig) EIP 2ACE133E      XCS 00000023
EFL 00000207      ESP 7FFFF654      XSS 0000002B
```

Stack: 1740 bytes. ESP 7FFFF654, TOP 7FFFFD20

```
0x7FFFF654: 00000000 00000008 00000003 08051E95 .....
0x7FFFF664: 00000005 7FFFF8CC 00000000 00000000 .....
0x7FFFF674: 7FFFF6CC 00000001 7FFFF95C 080522CD .....\"..
0x7FFFF684: 7FFFF9A4 00000008 7FFFFC34 2AC1F18C .....4.....*
```

Example 41-5 Displays All Process Log Details

```
switch# show processes log details
=====
Service: snmpd
Description: SNMP Agent

Started at Wed Jan  9 00:14:55 1980 (597263 us)
Stopped at Fri Jan 11 10:08:36 1980 (649860 us)
Uptime: 2 days 9 hours 53 minutes 53 seconds

Start type: SRV_OPTION_RESTART_STATEFUL (24)
Death reason: SYSMGR_DEATH_REASON_FAILURE_SIGNAL (2)
Exit code: signal 6 (core dumped)
CWD: /var/sysmgr/work

Virtual Memory:

CODE      08048000 - 0804C4A0
DATA      0804D4A0 - 0804D770
BRK       0804DFC4 - 0818F000
STACK     7FFFFCE0
TOTAL     26656 KB
...
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 41-6 Displays Memory Information About Processes

```
switch# show processes memory
PID      MemAlloc  StackBase/Ptr      Process
-----
1277     120632   7ffffcd0/7ffffefe4 xbar_client
1278       56800   7ffffce0/7ffffb5c  wwn
1279    1210220   7ffffce0/7ffffbac  vsan
1293     386144   7ffffcf0/7ffffebd4 span
1294    1396892   7ffffce0/7ffffdff4 snmpd
1295     214528   7ffffcf0/7ffff904  rscn
1296      42064   7ffffce0/7ffffb5c  qos
```

Where:

- MemAlloc = total memory allocated by the process.
- StackBase/Ptr = process stack base and current stack pointer in hex format.

Displaying System Status

Use the **show system** command to display system-related status information (see [Example 41-7](#) to [Example 41-10](#)).

Example 41-7 Displays Default Switch Port States

```
switch# show system default switchport
System default port state is down
System default trunk mode is on
```

Example 41-8 Displays Error Information for a Specified ID

```
switch# show system error-id 0x401D0019
Error Facility: module
Error Description: Failed to stop Linecard Async Notification.
```

Example 41-9 Displays the System Reset Information

```
switch# Show system reset-reason module 5
----- reset reason for module 5 -----
1) At 224801 usecs after Fri Nov 21 16:36:40 2003
   Reason: Reset Requested by CLI command reload
   Service:
   Version: 1.3(1)
2) At 922828 usecs after Fri Nov 21 16:02:48 2003
   Reason: Reset Requested by CLI command reload
   Service:
   Version: 1.3(1)
3) At 318034 usecs after Fri Nov 21 14:03:36 2003
   Reason: Reset Requested by CLI command reload
   Service:
   Version: 1.3(1)
4) At 255842 usecs after Wed Nov 19 00:07:49 2003
   Reason: Reset Requested by CLI command reload
   Service:
   Version: 1.3(1)
```

Send documentation comments to mdsfeedback-doc@cisco.com.

The **show system reset-reason** command displays the following information:

- In a Cisco MDS 9500 Series switch, the last four reset-reason codes for the supervisor module in slot 5 and slot 6 are displayed. If either supervisor module is absent, the reset-reason codes for that supervisor module are not displayed.
- In a Cisco MDS 9200 Series switch, the last four reset-reason codes for the supervisor module in slot 1 are displayed.
- The **show system reset-reason module *number*** command displays the last four reset-reason codes for a specific module in a given slot. If a module is absent, then the reset-reason codes for that module are not displayed.

Use the **clear system reset-reason** command to clear the reset-reason information stored in NVRAM and volatile persistent storage.

- In a Cisco MDS 9500 Series switch, this command clears the reset-reason information stored in NVRAM and volatile persistent storage in the active and standby supervisor modules.
- In a Cisco MDS 9200 Series switch, this command clears the reset-reason information stored in NVRAM and volatile persistent storage in the active supervisor module.

Example 41-10 Displays System Uptime

```
switch# show system uptime
Start Time: Sun Oct 13 18:09:23 2030
Up Time:    0 days, 9 hours, 46 minutes, 26 seconds
```

Use the **show system resources** command to display system-related CPU and memory statistics (see [Example 41-11](#)).

Example 41-11 Displays System-Related CPU and Memory Information

```
switch# show system resources
Load average:  1 minute: 0.43   5 minutes: 0.17   15 minutes: 0.11
Processes   :  100 total, 2 running
CPU states  :  0.0% user,   0.0% kernel,  100.0% idle
Memory usage: 1027628K total,   313424K used,   714204K free
                  3620K buffers,   22278K cache
```

Where:

- Load average—Displays the number of running processes. The average reflects the system load over the past 1, 5, and 15 minutes.
- Processes—Displays the number of processes in the system, and how many are actually running when the command is issued.
- CPU states—Displays the CPU usage percentage in user mode, kernel mode, and idle time in the last one second.
- Memory usage—Displays the total memory, used memory, free memory, memory used for buffers, and memory used for cache in KB. Buffers and cache are also included in the *used* memory statistics.

Send documentation comments to mdsfeedback-doc@cisco.com.

Core and Log Files

You can save cores (from the active supervisor module, the standby supervisor module, or any switching module) to an external Flash (slot 0) or to a TFTP server in one of two ways:

- On demand—Copies a single file based on the provided process ID.
- Periodically—Copies core files periodically as configured by the user.

A new scheme overwrites any previously issued scheme. For example, if you perform another core log copy task, the cores are periodically saved to the new location or file.



Tip

Be sure to create any required directory before performing this task. If the directory specified by this task does not exist, the switch software logs a system message each time a copy cores is attempted.

To copy the core and log files on demand, follow this step:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# copy core:7407 slot0:coreSample | Copies the core file with the process ID 7407 as coreSample in slot 0. |
| | switch# copy core://5/1524 tftp://1.1.1.1/abcd | Copies cores (if any) of a process with PID 1524 generated on slot 5 to a TFTP server. |

- If the core file for the specified process ID is not available, you see the following response:

```
switch# copy core:133 slot0:foo
No core file found with pid 133
```

- If two core files exist with the same process ID, only one file is copied:

```
switch# copy core:7407 slot0:foo1
2 core files found with pid 7407
Only "/isan/tmp/logs/calc_server_log.7407.tar.gz" will be copied to the destination.
```

To copy the core and log files periodically, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# system cores slot0:coreSample | Copies the core file coreSample to slot 0. |
| | switch(config)# system cores tftp://1.1.1.1/abcd | Copies the core file (abcd) in the specified directory on the TFTP server. |
| | switch(config)# no system cores | Disables the core files copying feature. |

Saving the Last Core to Flash

Prior to Cisco SAN-OS Release 2.0(1b), the last core dump(service core) is lost when one of the following events occur:

- A supervisor switchover in a Cisco MDS 9500 Series director
- A reboot in a single-supervisor Cisco MDS switch or in any switch in the Cisco MDS 9100 Series.

This last core dump that triggers the reset or switchover is saved in the RAM which is cleaned up at the end of the switchover or reboot process.

Send documentation comments to mdsfeedback-doc@cisco.com.

As of Cisco SAN-OS Release 2.0(1b), this last core dump is automatically saved to the flash in the /mnt/pss/ partition before the switchover or reboot occurs. Three minutes after the supervisor module reboots, the saved last core is restored from the flash partition (/mnt/pss) back its original RAM location. This restoration is a background process and is not visible to the user.



Tip

The timestamp on the restored last core file displays the time when the supervisor booted up—not when the last core was actually dumped. To obtain the exact time of the last core dump, check the corresponding log file with the same PID.

To view the last core information, issue the **show cores** command in EXEC mode.

To view the time of the actual last core dump, issue the **show process log** command in EXEC mode.

Clearing the Core Directory

Use the **clear cores** command to clean out the core directory. The software keeps the last few cores per service and per slot and clears all other cores present on the active supervisor module.

```
switch# clear cores
```

Displaying Core Status

Use the **show system cores** command to display the currently configured scheme for copying cores. See Examples 41-12 to 41-14.

Example 41-12 Displays the Status of System Cores

```
switch# show system cores
Transfer of cores is enabled
```

Example 41-13 Displays All Cores Available for Upload from the Active Supervisor Module

```
switch# show cores
Module-num Process-name  PID      Core-create-time
-----
5          fspf             1524     Nov 9 03:11
6          fcc              919      Nov 9 03:09
8          acltcam          285      Nov 9 03:09
8          fib              283      Nov 9 03:08
```

Where:

Module-num shows the slot number on which the core was generated. In this example, the fspf core was generated on the active supervisor module (slot 5), fcc was generated on the standby supervisor module (slot 6), and acltcam and fib were generated on the switching module (slot 8).

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 41-14 Displays Logs on the Local System

```
switch# show processes log
```

| Process | PID | Normal-exit | Stack | Core | Log-create-time |
|---------------|-------|-------------|-------|------|--------------------------|
| ExceptionLog | 2862 | N | Y | N | Wed Aug 6 15:08:34 2003 |
| acl | 2299 | N | Y | N | Tue Oct 28 02:50:01 2003 |
| bios_daemon | 2227 | N | Y | N | Mon Sep 29 15:30:51 2003 |
| capability | 2373 | N | Y | N | Tue Aug 19 13:30:02 2003 |
| core-client | 2262 | N | Y | N | Mon Sep 29 15:30:51 2003 |
| fcanalyzer | 5623 | N | Y | N | Fri Sep 26 20:45:09 2003 |
| fcd | 12996 | N | Y | N | Fri Oct 17 20:35:01 2003 |
| fcdomain | 2410 | N | Y | N | Thu Jun 12 09:30:58 2003 |
| ficon | 2708 | N | Y | N | Wed Nov 12 18:34:02 2003 |
| ficonstat | 9640 | N | Y | N | Tue Sep 30 22:55:03 2003 |
| flogi | 1300 | N | Y | N | Fri Jun 20 08:52:33 2003 |
| idehsd | 2176 | N | Y | N | Tue Jun 24 05:10:56 2003 |
| lmgrd | 2220 | N | N | N | Mon Sep 29 15:30:51 2003 |
| platform | 2840 | N | Y | N | Sat Oct 11 18:29:42 2003 |
| port-security | 3098 | N | Y | N | Sun Sep 14 22:10:28 2003 |
| port | 11818 | N | Y | N | Mon Nov 17 23:13:37 2003 |
| rlir | 3195 | N | Y | N | Fri Jun 27 18:01:05 2003 |
| rscn | 2319 | N | Y | N | Mon Sep 29 21:19:14 2003 |
| securityd | 2239 | N | N | N | Thu Oct 16 18:51:39 2003 |
| snmpd | 2364 | N | Y | N | Mon Nov 17 23:19:39 2003 |
| span | 2220 | N | Y | N | Mon Sep 29 21:19:13 2003 |
| syslogd | 2076 | N | Y | N | Sat Oct 11 18:29:40 2003 |
| tcap | 2864 | N | Y | N | Wed Aug 6 15:09:04 2003 |
| tftpd | 2021 | N | Y | N | Mon Sep 29 15:30:51 2003 |
| vpm | 2930 | N | N | N | Mon Nov 17 19:14:33 2003 |

Kernel Core Dumps



Caution

Changes to the kernel cores should be made by an administrator or individual who is completely familiar with switch operations.

When a specific module's operating system (OS) crashes, it is sometimes useful to obtain a full copy of the memory image (called a kernel core dump) to identify the cause of the crash. When the module experiences a kernel core dump it triggers the proxy server configured on the supervisor. The supervisor sends the module's OS kernel core dump to the Cisco MDS 9000 System Debug Server. Similarly, if the supervisor OS fails, the supervisor sends its OS kernel core dump to the Cisco MDS 9000 System Debug Server.



Note

The Cisco MDS 9000 System Debug Server is a Cisco application that runs on Linux. It creates a repository for kernel core dumps. You can download the Cisco MDS 9000 System Debug Server from the Cisco.com website at <http://www.cisco.com/kobayashi/sw-center/sw-stornet.shtml>.

Kernel core dumps are only useful to your technical support representative. The kernel core dump file, which is a large binary file, must be transferred to an external server that resides on the same physical LAN as the switch. The core dump is subsequently interpreted by technical personnel who have access to source code and detailed memory maps.

Send documentation comments to mdsfeedback-doc@cisco.com.



Tip

Core dumps take up disk space on the Cisco MDS 9000 System Debug Server application. If all levels of core dumps (**level all** option) are configured, you need to ensure that a minimum of 1 GB of disk space is available on the Linux server running the Cisco MDS 9000 System Debug Server application to accept the dump. If the process does not have sufficient space to complete the generation, the module resets itself. All changes made to kernel cores are saved to the running configuration.

To configure the external server, follow these steps:

| | Command | Purpose |
|---------------|--|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# kernel core target 10.50.5.5 succeeded | Configures the external server's IP address. |

To configure the module information, follow these steps:

| | Command | Purpose |
|---------------|---|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# kernel core module 5 succeeded | Configures kernel core generation for module 5. |
| | switch(config)# kernel core module 5 level header succeeded | Configures kernel core generation for module 5, and limits the generation to header-level cores. |
| Step 3 | switch(config)# kernel core limit 2 succeeded | Configures kernel core generations for two modules. The default is 1 module. |

All changes made to kernel cores may be viewed using the **show running-config** command. Alternatively, use the **show kernel cores** command to view specific configuration changes (see [Example 41-15](#) to [Example 41-17](#)).

Example 41-15 Displays the Core Limit

```
switch# show kernel core limit
2
```

Example 41-16 Displays the External Server

```
switch# show kernel core target
10.50.5.5
```

Example 41-17 Displays the Core Settings for the Specified Module

```
switch# show kernel core module 5
module 5 core is enabled
  level is header
  dst_ip is 10.50.5.5
  src_port is 6671
  dst_port is 6666
  dump_dev_name is eth1
  dst_mac_addr is 00:00:0C:07:AC:01
```

Send documentation comments to mdsfeedback-doc@cisco.com.

Online System Health Management

The Online Health Management System (system health) is a hardware fault detection and recovery feature. It ensures the general health of switching, services, and supervisor modules in any switch in the Cisco MDS 9000 Family as of Cisco MDS SAN-OS Release 1.3(4) and later.

The system health application runs on all Cisco MDS modules and monitors system hardware in a given MDS switch. The system health application running in the standby supervisor module only monitors the standby supervisor module—if that module is available in the HA standby mode.

See the “HA Switchover Characteristics” section on page 5-2.

The system health application launches a daemon process in all modules and runs multiple tests on each module to test individual module components. The tests run at pre-configured intervals, cover all major fault points, and isolate any failing component in the MDS switch. The system health running on the active supervisor maintains control over all other system health components running on all other modules in the switch.

On detecting a fault, the system health application attempts the following recovery actions:

- Sends Call Home and system messages and exception logs as soon as it detects a failure.
- Shuts down the failing module or component (such as an interface).
- Isolates failed ports from further testing.
- Reports the failure to the appropriate software component.
- Switches to the standby supervisor module, if an error is detected on the active supervisor module and a standby supervisor module exists in the Cisco MDS switch. After the switchover, the new active supervisor module restarts the active supervisor tests.
- Reloads the switch if a standby supervisor module does not exist in the switch.
- Provides CLI support to view, test, and obtain test run statistics or change the system health test configuration on the switch.
- Performs tests to focus on the problem area:
- Retrieves its configuration information from persistent storage.

Each module is configured to run the test relevant to that module. You can change the default parameters of the test in each module as required.

Send documentation comments to mdsfeedback-doc@cisco.com.

System Health Initiation

By default, the system health feature is enabled in each switch in the Cisco MDS 9000 Family.

To disable or enable this feature in any switch in the Cisco MDS 9000 Family, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# no system health System Health is disabled. | Disables system health from running tests in this switch. |
| | switch(config)# system health System Health is enabled. | Enables (default) system health to run tests in this switch. |
| Step 3 | switch(config)# no system health interface fc8/1 System health for interface fc8/13 is disabled. | Disables system health from testing the specified interface. |
| | switch(config)# system health interface fc8/1 System health for interface fc8/13 is enabled. | Enables (default) system health to test for the specified interface. |

Loopback Test Configuration Frequency

Loopback tests are designed to identify hardware errors in the data path in the module(s) and the control path in the supervisors. One loopback frame is sent to each module at a preconfigured frequency—it passes through each configured interface and returns to the supervisor module.

The loopback tests can be run at frequencies ranging from 5 seconds (default) to 255 seconds. If you do not configure the loopback frequency value, the default frequency of 5 seconds is used for all modules in the switch. Loopback test frequencies cannot be altered for each module. The configured value is constant for all modules.

To configure the frequency of loopback tests for all modules in any switch in the Cisco MDS 9000 Family, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# system health loopback frequency 50 The new frequency is set at 50 Seconds. | Configures the loopback frequency to 50 seconds. The default loopback frequency is 5 seconds. The valid range is from 5 to 255 seconds. |

Hardware Failure Action

The failure-action command controls the Cisco SAN-OS software from taking any action if a hardware failure is determined while running the tests.

By default, this feature is enabled in all switches in the Cisco MDS 9000 Family—action is taken if a failure is determined and the failed component is isolated from further testing.

Failure action is controlled at individual test levels (per module), at the module level (for all tests), or for the entire switch.

Send documentation comments to mdsfeedback-doc@cisco.com.

To configure failure action in a switch, follow these steps:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# system health failure-action System health global failure action is now enabled. | Enables the switch to take failure action (default) . |
| Step 3 | switch(config)# no system health failure-action System health global failure action now disabled. | Reverts the switch configuration to prevent failure action being taken. |
| Step 4 | switch(config)# system health module 1 failure-action System health failure action for module 1 is now enabled. | Enables switch to take failure action for failures in module 1. |
| Step 5 | switch(config)# no system health module 1 loopback failure-action System health failure action for module 1 loopback test is now disabled. | Prevents the switch from taking action on failures determined by the loopback test in module 1. |

Test Run Requirements

Enabling a test does not guarantee that a test will run.

Tests on a given interface or module only run if you enable system health for all of the following items:

- The entire switch.
- The required module.
- The required interface.



Tip

The test will not run if system health is disabled in any combination. If system health is disabled to run tests, the test status shows up as disabled.



Tip

If the specific module or interface is enabled to run tests, but is not running the tests due to system health being disabled, then tests shows up as enabled, (not running).

Tests for a Specified Module

The system health feature in the SAN-OS software performs tests in the following areas:

- Active supervisor's in-band connectivity to the fabric.
- Standby supervisor's arbiter availability.
- Boot flash connectivity and accessibility on all modules.
- EOBC connectivity and accessibility on all modules.
- Data path integrity for each interface on all modules.
- Management port's connectivity.
- Caching services module batteries (for temperature, age, full-charge capacity, (dis)charge ability and backup capability) and cache disks (for connectivity, accessibility and raw disk I/O).

Send documentation comments to mdsfeedback-doc@cisco.com.

- User-driven test for external connectivity verification, port is shutdown during the test (FC ports only).
- User-driven test for internal connectivity verification (Fibre Channel and iSCSI ports).

To perform the required test on a specific module, follow these steps:

| | Command | Purpose |
|---------------|---|---|
| Step 1 | switch# config terminal switch(config)# | Enters configuration mode. |
| Note | The following steps can be performed in any order. | |
| Step 2 | switch(config)# system health module 8 battery-charger battery-charger test is not configured to run on module 8. | Enables the battery-charger test on both batteries in the CSM module residing in slot 8. If the switch does not have a CSM in slot 8, this message is issued, |
| Step 3 | switch(config)# system health module 8 cache-disk cache-disk test is not configured to run on module 8. | Enables the cache-disk test on both disks in the CSM module residing in slot 8. If the switch does not have a CSM in slot 8, this message is issued, |
| Note | The various options for each test are described in the next step. Each command can be configured in any order. The various options are presented in the same step for documentation purposes. | |
| Step 4 | switch(config)# system health module 8 bootflash System health for module 8 Bootflash is already enabled. | Enables the bootflash test on Module 8. |
| | switch(config)# system health module 8 bootflash frequency 200 The new frequency is set at 200 Seconds. | Sets the new frequency of the bootflash test on module 8 to 200 seconds. |
| Step 5 | switch(config)# system health module 8 eobc System health for module 8 EOBC is now enabled. | Enables the EOBC test on Module 8. |
| Step 6 | switch(config)# system health module 8 loopback System health for module 8 EOBC is now enabled. | Enables the loopback test on Module 8. |
| Step 7 | switch(config)# system health module 5 management System health for module 8 EOBC is now enabled. | Enables the management test on Module 5. |

Clearing Previous Error Reports

You can clear the error history for Fibre Channel interfaces, iSCSI interfaces, for an entire module, or one particular test for an entire module. By clearing the history, you are directing the software to retest all failed components that were previously excluded from tests.

If you previously enabled the **failure-action** option for a period of time (for example, one week) to prevent OHMS from taking any action when a failure is encountered and after that week you are now ready to start receiving these errors again, the you must clear the system health error status for each test.



Tip

The management port test cannot be run on a standby supervisor module.

Use the EXEC-level **system health clear-errors** command at the interface or module level to erase any previous error conditions logged by the system health application. The **battery-charger**, the **bootflash**, the **cache-disk**, the **eobc**, the **inband**, the **loopback**, and the **mgmt** test options can be individually specified for a given module.

Send documentation comments to mdsfeedback-doc@cisco.com.

The following example clears the error history for the specified Fibre Channel interface:

```
switch# system health clear-errors interface fc 3/1
```

The following example clears the error history for the specified module:

```
switch# system health clear-errors module 3
```

The following example clears the management port test error history for the specified module:

```
switch# system health clear-errors module 1 mgmt
```

Performing Internal Loopbacks

Internal loopback tests send and receive FC2 frames to/from the same ports and provides the round trip time taken in microseconds. These tests are available for both Fibre Channel and iSCSI interfaces.

Use the EXEC-level **system health internal-loopback** command to explicitly run this test on demand (when requested by the user) within ports for the entire module.

```
switch# system health internal-loopback interface iscsi 8/1
Internal loopback test on interface iscsi8/1 was successful.
Round trip time taken is 79 useconds
```



Note

If the test fails to complete successfully, the software analyzes the failure and prints the following error:
External loopback test on interface fc 7/2 failed. Failure reason: Failed to loopback, analysis complete Failed device ID 3 on module 1

Performing External Loopbacks

External loopback tests send and receive FC2 frames to/from the same port. You need to connect a cable (or a plug) to loop the Rx port to the Tx port before running the test. This test is only available for Fibre Channel interfaces.

Use the EXEC-level **system health external-loopback** command to run this test on demand for external devices connected to a switch that is part of a long-haul network.

```
switch# system health external-loopback interface fc 3/1
This will shut the requested interfaces Do you want to continue (y/n)? [n] y
External loopback test on interface fc3/1 was successful.
```

Use the **system health external-loopback interface force** command to shut down the required interface directly without a back out confirmation.

```
switch# system health external-loopback interface fc 3/1 force
External loopback test on interface fc3/1 was successful.
```



Note

If the test fails to complete successfully, the software analyzes the failure and prints the following error:
External loopback test on interface fc 7/2 failed. Failure reason: Failed to loopback, analysis complete Failed device ID 3 on module 1

Send documentation comments to mdsfeedback-doc@cisco.com.

Interpreting the Current Status

The status of each module or test depends on the current configured state of the OHMS test in that particular module (see [Table 41-1](#)).

Table 41-1 *OHMS Configured Status for Tests and Modules*

| Status | Description |
|------------------|---|
| Enabled | You have currently enabled the test in this module and the test is not running. |
| Disabled | You have currently disabled the test in this module. |
| Running | You have enabled the test and the test is currently running in this module. |
| Failing | This state is displayed if a failure is imminent for the test running in this module—possibility of test recovery exists in this state. |
| Failed | The test has failed in this module—and the state cannot be recovered. |
| Stopped | The test has been internally stopped in this module by the Cisco SAN-OS software. |
| Internal failure | The test encountered an internal failure in this module. For example, the system health application is not able to open a socket as part of the test procedure. |
| Diags failed | The startup diagnostics has failed for this module or interface. |
| On demand | The system health external-loopback or the system health internal-loopback tests are currently running in this module. Only these two commands can be issued on demand. |
| Suspended | Only encountered in the MDS 9100 Series due to one oversubscribed port moving to a E or TE port mode. If one oversubscribed port moves to this mode, the other three oversubscribed ports in the group are suspended. |

The status of each test in each module is visible when you display any of the **show system health** commands. See the “[Displaying System Health](#)” section on page 41-15.

Displaying System Health

Use the **show system health** command to display system-related status information (see [Example 41-18](#) to [Example 41-23](#)).

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 41-18 Displays the Current Health of All Modules in the Switch

```
switch# show system health
```

Current health information for module 2.

| Test | Frequency | Status | Action |
|-----------|-----------|---------|---------|
| ----- | ----- | ----- | ----- |
| Bootflash | 5 Sec | Running | Enabled |
| EOBC | 5 Sec | Running | Enabled |
| Loopback | 5 Sec | Running | Enabled |
| ----- | ----- | ----- | ----- |

Current health information for module 6.

| Test | Frequency | Status | Action |
|-----------------|-----------|---------|---------|
| ----- | ----- | ----- | ----- |
| InBand | 5 Sec | Running | Enabled |
| Bootflash | 5 Sec | Running | Enabled |
| EOBC | 5 Sec | Running | Enabled |
| Management Port | 5 Sec | Running | Enabled |
| ----- | ----- | ----- | ----- |

Example 41-19 Displays the Current Health of a Specified Module

```
switch# show system health module 8
```

Current health information for module 8.

| Test | Frequency | Status | Action |
|-----------|-----------|---------|---------|
| ----- | ----- | ----- | ----- |
| Bootflash | 5 Sec | Running | Enabled |
| EOBC | 5 Sec | Running | Enabled |
| Loopback | 5 Sec | Running | Enabled |
| ----- | ----- | ----- | ----- |

Example 41-20 Displays Health Statistics for All Modules

```
switch# show system health statistics
```

Test statistics for module # 1

| Test Name | State | Freq(s) | Run | Pass | Fail | CFail | Errs |
|-----------|---------|---------|-------|-------|-------|-------|-------|
| ----- | ----- | ----- | ----- | ----- | ----- | ----- | ----- |
| Bootflash | Running | 5s | 12900 | 12900 | 0 | 0 | 0 |
| EOBC | Running | 5s | 12900 | 12900 | 0 | 0 | 0 |
| Loopback | Running | 5s | 12900 | 12900 | 0 | 0 | 0 |
| ----- | ----- | ----- | ----- | ----- | ----- | ----- | ----- |

Test statistics for module # 3

| Test Name | State | Freq(s) | Run | Pass | Fail | CFail | Errs |
|-----------|---------|---------|-------|-------|-------|-------|-------|
| ----- | ----- | ----- | ----- | ----- | ----- | ----- | ----- |
| Bootflash | Running | 5s | 12890 | 12890 | 0 | 0 | 0 |
| EOBC | Running | 5s | 12890 | 12890 | 0 | 0 | 0 |
| Loopback | Running | 5s | 12892 | 12892 | 0 | 0 | 0 |
| ----- | ----- | ----- | ----- | ----- | ----- | ----- | ----- |

Test statistics for module # 5

| Test Name | State | Freq(s) | Run | Pass | Fail | CFail | Errs |
|-----------|-------|---------|-------|-------|-------|-------|-------|
| ----- | ----- | ----- | ----- | ----- | ----- | ----- | ----- |

Send documentation comments to mdsfeedback-doc@cisco.com.

```
InBand           Running           5s   12911   12911       0     0     0
Bootflash        Running           5s   12911   12911       0     0     0
EOBC             Running           5s   12911   12911       0     0     0
Management Port  Running           5s   12911   12911       0     0     0
-----
```

Test statistics for module # 6

```
-----
Test Name      State      Freq(s)   Run    Pass    Fail CFail Errs
-----
InBand         Running    5s        12907  12907    0     0     0
Bootflash      Running    5s        12907  12907    0     0     0
EOBC           Running    5s        12907  12907    0     0     0
-----
```

Test statistics for module # 8

```
-----
Test Name      State      Freq(s)   Run    Pass    Fail CFail Errs
-----
Bootflash      Running    5s        12895  12895    0     0     0
EOBC           Running    5s        12895  12895    0     0     0
Loopback       Running    5s        12896  12896    0     0     0
-----
```

Example 41-21 Displays Statistics for a Specified Module

```
switch# show system health statistics module 3
```

Test statistics for module # 3

```
-----
Test Name      State      Freq(s)   Run    Pass    Fail CFail Errs
-----
Bootflash      Running    5s        12932  12932    0     0     0
EOBC           Running    5s        12932  12932    0     0     0
Loopback       Running    5s        12934  12934    0     0     0
-----
```

Example 41-22 Displays Loopback Test Statistics for the Entire Switch

```
switch# show system health statistics loopback
```

```
-----
Mod Port Status      Run    Pass    Fail    CFail Errs
-----
1 16 Running    12953  12953    0        0     0
3 32 Running    12945  12945    0        0     0
8 8 Running     12949  12949    0        0     0
-----
```

Example 41-23 Displays Loopback Test Statistics for a Specified Interface

```
switch# show system health statistics loopback interface fc 3/1
```

```
-----
Mod Port Status      Run    Pass    Fail    CFail Errs
-----
3 1 Running         0        0        0        0     0
-----
```



Note

Interface-specific counters will remain at zero unless the module-specific loopback test reports errors or failures.

Send documentation comments to mdsfeedback-doc@cisco.com.

Example 41-24 Displays the Loopback Test Time Log for All Modules

```
switch# show system health statistics loopback timelog
```

| Mod | Samples | Min (usecs) | Max (usecs) | Ave (usecs) |
|-----|---------|-------------|-------------|-------------|
| 1 | 1872 | 149 | 364 | 222 |
| 3 | 1862 | 415 | 743 | 549 |
| 8 | 1865 | 134 | 455 | 349 |

Example 41-25 Displays the Loopback Test Time Log for a Specified Module

```
switch# show system health statistics loopback module 8 timelog
```

| Mod | Samples | Min (usecs) | Max (usecs) | Ave (usecs) |
|-----|---------|-------------|-------------|-------------|
| 8 | 1867 | 134 | 455 | 349 |

Default Settings

Table 41-2 lists the default system health and log settings.

Table 41-2 Default System Health and Log Settings

| Parameters | Default |
|------------------------|------------|
| Kernel core generation | One module |
| System health | Enabled |
| Loopback frequency | 5 seconds |
| Failure action | Enabled. |



Symbols

* (asterisk)

autolearned entries [21-11](#)

host time stamp [27-46](#)

IKE version 1 tunnel [29-33](#)

iSCSI node [28-114](#)

port security wildcard [21-9](#)

Numerics

14/2-port Multiprotocol Services module. See MPS-14/2 module

16-port switching modules

asset tags [12-17](#)

configuring BB_credits [12-12](#)

LEDs [12-17](#)

port groups [12-17](#)

See also switching modules

32-port switching modules

configuration guidelines [12-8](#)

configuring BB_credits [12-12](#)

PortChannel configuration guidelines [14-3](#)

SPAN guidelines [38-7](#)

See also switching modules

A

AAA

authorization and authentication process [19-20](#)

setting authentication [19-19](#)

usage [1-13](#), [19-1](#)

Access Control Lists. See ACLs

accounting [19-32](#)

ACL based access control

configuring for iSCSI [28-69](#)

ACLs

adding entries [26-9](#)

applying [26-10](#)

clearing counters [26-11](#)

configuration guidelines [26-5](#)

creating [26-8](#)

crypto [29-12 to 29-16](#)

defining [26-8](#)

operands [26-8](#)

reading log dumps [26-9](#)

removing entries [26-9](#)

activation

fabric binding [27-39](#)

active equals saved command [27-19](#), [27-25](#)

active zone sets

considerations [15-8](#)

distributing [15-12](#)

adding

IP addresses [26-21](#)

adding ACL entries [26-9](#)

address-allocation cache [31-16](#)

Address Resolution Protocol. See ARP

administrative speed

configuring [12-11](#)

administrative states

description [12-6](#)

administrator passwords

default [4-6](#)

recovering [19-37](#)

requirements (note) [4-6](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

advertisement packets

setting time intervals [26-22](#)

AF IDs

description [17-10](#)

aggregated flow statistics [24-15](#)

aliases

configuring [15-6](#)

ALPA allocation [12-30](#)

ALPA caches

clearing [12-11](#)

configuring [12-10](#)

displaying contents [12-11](#)

area FCID

configuring [31-11](#)

ARP

clearing and viewing entries [26-13](#)

IP services [1-8](#)

ARP caches

clearing [28-11](#)

displaying [28-11](#)

ASMs

reloading [7-20](#)

specifying image boot variables [7-15](#)

specifying images [7-19](#)

ASM-SFN boot variables

configuring ASMs for VSFN [7-15](#)

configuring SSMs for VSFN [7-19](#)

ASM-SFN variables

specifying [7-20](#)

asset tags

16-port switching modules [12-17](#)

assigning

domain IDs [31-4](#)

FC IDs [24-9](#)

global keys [19-6](#)

host key [19-6](#)

authentication

CHAP option [28-94](#)

iSCSI setup [28-93](#)

See also MD5 authentication

See also simple text authentication

authentication, authorization, and accounting. See AAA

automatic synchronization

conditions [5-5](#)

auto mode

configuring [12-10](#)

description [12-5](#)

auto-negotiation

configuring Gigabit Ethernet interfaces [28-6](#)

autonomous fabric IDs. see AF IDs [17-10](#)

AutoNotify

destination profile (note) [30-4](#)

registration [30-2](#)

service contract [30-3](#)

auto port mode

interface configuration [12-3](#)

B

basic input/output system. See BIOS

BB_credits

configuring [12-12](#)

port swapping [27-26](#)

reason codes [12-7](#)

beacon modes

configuring [12-16](#)

identifying LEDs [12-17](#)

Berkeley Packet Filter. See BPF

BIOS

boot sequence [6-25](#)

recovering corrupted bootflash [6-27](#)

recovery sequence [6-26](#)

setup (figure) [6-28](#)

BIOS upgrades [6-21](#)

blocking ports [27-27](#)

boot

sequence [6-25](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

bootflash

- copying to [6-18](#)
- device [6-18](#)
- file system [6-2](#)
- recovering corrupted [6-25 to 6-26](#)
- space requirements [6-2](#)

bootflash:. See internal bootflash:

bootloader

- nondisruptive upgrades [6-19](#)
- skipping phases [6-30](#)

boot variables

- configuring ASM-SFN boot variables [7-15, 7-19](#)
- configuring automatic copying to standby supervisor modules [5-4](#)
- configuring SSI boot variables [7-17, 7-21](#)
- disruptive upgrades [6-23](#)
- specifying ASM boot images [7-15 to 7-18](#)
- synchronizing [5-4](#)

BPF

- library [39-17](#)
- See also libpcap freeware

B port mode

- description [12-5](#)
- interface modes [12-5](#)

B ports

- configuring [28-36](#)
- interoperability mode [28-34](#)
- SAN extenders [28-35](#)

bridge port mode. See B port mode

bridge ports. See B ports

broadcast

- in-band addresses default [7-23](#)
- routing [24-10, 24-11](#)

buffer sizes

- configuring in FCIP profiles [28-29](#)

buffer-to-buffer credits

- See BB_credits

build fabric frames

- description [31-3](#)

C

cache. See address-allocation cache

Call Home

- configuring [30-3 to 30-14](#)
- description [1-10](#)
- message format options [30-2](#)
- periodic inventory notification [30-11](#)

capture filters [39-17](#)

CDP

- clearing [4-41](#)
- configuring [4-40 to 4-45](#)
- configuring hold time [4-41](#)
- configuring refresh time interval globally [4-41](#)
- configuring version [4-41](#)
- disabling globally [4-40](#)
- disabling on an interface [4-40](#)
- displaying information [4-42](#)
- packet transmission [4-40](#)

CFS

- application requirements [9-5](#)
- clearing session locks [9-7](#)
- committing changes [9-6](#)
- default settings [9-11](#)
- description [9-2](#)
- disabling on a switch [9-5](#)
- discarding changes [9-6](#)
- displaying information [9-7](#)
- distribution modes [9-4](#)
- distribution scopes [9-4](#)
- enabling [9-5](#)
- fabric locking [9-6](#)
- feature description [9-3](#)
- merge support [9-7](#)
- protocol description [9-3](#)
- SAN-OS features supported [9-2](#)
- saving configurations [9-6](#)
- CHAP authentication [28-94](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

CIM

configuring [12-20](#)

CIM servers

configuring security [12-20](#)

displaying information [12-25](#)

Cisco Discovery Protocol. See CDP

Cisco Fabric Services. See CFS

Cisco MDS 9000 Family

initial setup [4-2 to 4-13](#)

starting a switch [4-2](#)

Cisco MDS 9000 System Debug Server [41-8](#)

Cisco MDS 9100 Series

high availability [1-5, 5-1](#)

overview [1-2](#)

Cisco MDS 9120

overview [1-2](#)

Cisco MDS 9140

overview [1-2](#)

Cisco MDS 9200 Series

high availability [1-5, 5-2](#)

overview [1-2](#)

supervisor modules [1-11](#)

Cisco MDS 9216A switches

overview [1-2](#)

Cisco MDS 9216i switches

bundled licenses [3-3](#)

configuring extended BB_credits [12-14](#)

overview [1-2](#)

Cisco MDS 9216 switches

high availability [5-1](#)

overview [1-2](#)

Cisco MDS 9500 Series

high availability [1-4, 5-1](#)

overview [1-3](#)

Cisco MDS 9506 Directors

overview [1-3](#)

supervisor modules [1-11](#)

Cisco MDS 9509 Directors

overview [1-3](#)

supervisor modules [1-11](#)

Cisco MDS SAN-OS

downgrading [4-31](#)

software images [6-1](#)

clearing

DPVM database [11-5](#)

FIB statistics [24-16](#)

FICON device allegiance [27-27](#)

FSPF counters [24-9](#)

CLI

accessing submodes [2-3](#)

command modes [2-3](#)

command prompt [2-2](#)

command scripts [2-27](#)

description [1-15](#)

Fabric Manager alternative [1-15](#)

setting delay time [2-28](#)

clock

mainframe [27-17](#)

clock modules

description [8-11](#)

displaying status [8-11](#)

CMOS

configuration [6-28](#)

saving changes [6-29](#)

code page

FICON options [27-16](#)

COM1 ports

configuring settings [4-35](#)

verifying settings [4-35](#)

command-line interface. See CLI

commands

saving output to files [2-26](#)

command scheduler

description [1-14](#)

overview [35-1](#)

Common Information Model. See CIM

Send documentation comments to mdsfeedback-doc@cisco.com.

CompactFlash

- devices [6-19](#)
- disk [6-2](#)
- slot 0 [6-19](#)

CompactFlash. See external CompactFlash

Company ID

- FC ID allocation [39-20](#)

computing routes [24-1](#)

conditional receive

- RLIR [27-46](#)

configuration

- backing up current [4-29](#)
- restoring redundancy mode [4-30](#)
- rolling back to previous [4-30](#)
- saving to NVRAM [4-27](#)

configuration files

- displaying [4-25](#)
- distributing to fabric [4-28](#)
- downloading [4-25](#)
- FICON [27-23](#)
- saving [4-26](#)

configuring

- unique area FCIDs [31-12](#)

congestion control methods. See FCC; edge quench congestion control

congestion window monitoring. See CWM

connecting a modem

- COM 1 [4-36](#)
- console [4-36](#)

console ports

- configuring settings [4-34](#)
- verifying settings [4-34](#)

console session

- severity levels [36-4](#)

contact information

- assigning [30-3](#)

control traffic

- disabling [32-4](#)

Control Unit Port. See CUP

core dumps

- IPS module [28-120](#)
- kernel [41-8](#)

cores [41-6](#)

CUP

- blocking restriction [27-21](#)
- in-band management [27-1, 27-27](#)

current directory

- displaying [2-23](#)
- setting [2-22](#)

customized

- targets [37-4](#)

CWM

- configuring in FCIP profiles [28-28](#)

D

data field sizes

- configuring [12-16](#)

date

- configuring [4-16](#)

dead time interval [24-6](#)

default gateway

- BIOS setup configuration [6-28](#)
- configuring [4-23](#)
- recovering loader> prompt [6-30](#)
- recovering switch(boot)# prompt [6-31](#)

default gateways

- configuring mgmt0 interfaces [12-19](#)

default user

- description [4-3](#)

default VSAN

- description [10-5](#)

default zones

- description [15-11](#)
- interoperability [39-23](#)

deleting

- FSPF configurations [24-5](#)

deny conditions [26-5, 26-8](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

destination IDs

- exchange based [14-6](#)
- flow based [14-5](#)
- frame identification [32-2](#)
- frame loop back [39-5](#)
- in-order delivery [24-11, 32-2](#)
- load balancing [1-8](#)
- path selection [10-6](#)

destination profiles

- configuring [30-4](#)

device alias database

- committing changes [16-4](#)
- discarding changes [16-4](#)
- locking the fabric [16-3](#)
- merging [16-5](#)

device aliases

- clearing statistics [16-6](#)
- comparison with zones (table) [16-2](#)
- default settings [16-10](#)
- description [16-1](#)
- displaying information [16-6 to 16-9](#)
- distribution to fabric [16-5](#)
- features [16-2](#)
- modifying the database [16-3](#)
- overriding fabric lock [16-4](#)
- requirements [16-2](#)
- zone alias conversion [16-5](#)

device allegiance

- FICON [27-27](#)

device IDs

- Call Home format [30-21, 30-22](#)
- report capacity [37-1](#)

Device Manager

- description [1-15](#)

differentiated services code point. See DSCP

digital signature algorithm. See DSA key pairs

Dijkstra's algorithm [24-2](#)

direct memory access. See DMA-bridge

directories

- creating [2-23](#)
- deleting [2-24](#)
- display current [2-23](#)

disabling routing protocols [24-5](#)

discovered

- LUNs [37-4](#)
- targets [37-3](#)

display filters

- selective viewing [39-13](#)

disruptive

- upgrades [6-4](#)

distribution tree [24-10](#)

DMA-bridge

- displaying statistics [28-13](#)

documentation

- additional publications [xliv](#)
- related documents [xliv](#)

domain ID

- IVR guidelines [17-4](#)

domain IDs

- assignment failures [12-7](#)
- configuring [31-4](#)
- configuring zone members [15-4](#)
- distributing [31-2](#)
- interoperability [39-23](#)
- preferred [31-5](#)
- static [31-5](#)
- unique [17-6](#)

domain manager

- isolation [12-7](#)

domain names

- defining [26-26](#)

Domain Name System servers. See DNS servers

domain overlap

- isolation [12-7](#)

downgrading

- Cisco MDS SAN-OS releases [4-31](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

DPVM

- default settings [11-12](#)
- description [1-6, 11-2](#)
- displaying configurations [11-8](#)
- enabling [11-2](#)
- requirements [11-2](#)
- sample configuration [11-10 to 11-12](#)

DPVM databases

- activating config databases [11-3](#)
- autolearned entries [11-4](#)
- comparing [11-8](#)
- configuring distribution [11-5 to 11-7](#)
- copying [11-7](#)
- description [11-3](#)
- displaying [11-9](#)
- enabling autolearning [11-4](#)
- merging guidelines [11-7](#)

drivers

- iSCSI [28-51](#)

drop latency time

- configuring [24-14](#)

dsa key pairs

- generating [19-35](#)

DSCP

- configuring [28-36](#)

duplicate messages

- throttling [30-12](#)

Dynamic Port VSAN Membership. See DPVM

dynamic VSANs. See DPVM

E

EBCDIC

- FICON string format [27-16](#)

edge quench congestion control

- description [32-2](#)

EFMD

- fabric binding [27-37](#)
- egress port [38-12, 38-25](#)

EISL

- trunking [1-8](#)

EISLs

- PortChannel links [14-2](#)

e-mail notification

- Call Home [30-1](#)

enforcing licenses [3-1](#)

Enhanced ISL. See EISL

enhanced zones

- advantages over basic zones [15-28](#)
- broadcast frames [15-32](#)
- changing from basic zones [15-28](#)
- configuring default policies [15-32](#)
- creating attribute [15-30](#)
- default settings [15-36](#)
- description [15-27](#)
- displaying information [15-33](#)
- enabling [15-29](#)
- merging databases [15-31](#)
- modifying database [15-30](#)

enterprise package licenses

- description [3-4](#)

EPLD images

- downgrading [7-14](#)
- upgrading [7-13](#)

E port mode

- classes of service [12-3](#)
- description [12-3](#)

E ports

- 32-port guidelines [12-8](#)
- 32-port switching module configuration guidelines [14-3](#)
- configuring [12-10, 28-36](#)
- FSPF topology [24-2](#)
- isolation [12-7](#)
- recovering from isolation [15-13](#)
- SPAN [38-4](#)
- trunking [1-8](#)
- trunking configuration [13-2](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

error messages

description [36-1](#)

error state [6-33](#)

Ethereal freeware

analyzer [39-10](#)

information [39-8](#)

Ethernet MAC statistics

displaying [28-12](#)

Ethernet PortChannel aggregation

description [28-17](#)

Ethernet PortChannels

adding Gigabit Ethernet interfaces [28-19](#)

redundancy [28-50](#)

Exchange Fabric Membership Data

see EFMD [27-37](#)

exchange IDs

in-order delivery [24-11](#)

load balancing [1-8, 39-5](#)

path selection [10-6](#)

exchange link parameter. See ELP

expansion port mode. See E port mode

expiry alerts

licenses [3-11](#)

extended BB_credits

configuring [xli, 12-14](#)

Extended Binary-Coded Decimal Interchange Code

see EBCDIC [27-16](#)

external CompactFlash

description [2-20](#)

formatting [2-21](#)

recovering from corruption [2-21](#)

supported devices [2-22](#)

external RADIUS server

CHAP [28-94](#)

external server

configuring [41-9](#)

F

fabric

See reconfigure fabric frames; build fabric frames

See build fabric frames

Fabric Analyzer

configuring [39-10](#)

description [39-8](#)

fabric binding

configuration [27-37](#)

default settings [27-49](#)

enforcement [27-38](#)

forceful activation [27-40](#)

port security comparison [27-37](#)

Fabric Configuration Server. See FCS

Fabric-Device Management Interface. See FDMI

fabric login. See FLOGI

fabric loop port mode. See FL port mode

Fabric Manager

description [1-12, 1-15](#)

Fabric Manager Server package license

description [3-5](#)

fabric names

setting [31-8](#)

fabric port mode. See F port mode

fabric pWWNs

configuring zone members [15-4](#)

zone membership [15-2](#)

fabric reconfiguration

fcdomain phase [31-2](#)

fabric security

default settings [20-12](#)

fan modules

description [8-10](#)

displaying status [8-11](#)

fault tolerant fabric

example (figure) [24-2](#)

FC aliases

configuring zone members [15-4](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

fcaliases

- cloning [15-21](#)
- renaming [15-21](#)

fcanalyzer

- clearing hosts [39-12](#)
- displaying filters [39-13](#)

FCC

- benefits [32-2](#)
- default settings [32-14](#)
- enabling [32-3](#)
- frame handling [32-2](#)
- logging facility [36-2](#)

fcdomain

- configuring [31-1](#)
- default settings [31-17](#)

FC ID allocation

- FICON implementation [27-10](#)

FC IDs

- allocating [31-2, 39-20](#)
- allocating Company IDs [39-20](#)
- configuring zone members [15-4](#)
- last byte [27-16](#)

FCIP

- configuring [28-19 to 28-42](#)
- configuring compression [28-41](#)
- configuring tape acceleration [28-39 to 28-41](#)
- configuring write acceleration [28-37 to 28-39](#)
- default parameters [28-121](#)
- discarding packets [28-32](#)
- displaying information [28-42](#)
- FICON support [27-4](#)
- Gigabit Ethernet ports [28-5](#)
- high availability [28-48 to 28-51](#)
- interfaces [28-22](#)
- IPS module [28-20](#)
- IPS module support [28-2](#)
- IP storage services support [28-2](#)
- link failures [28-49](#)
- MPS-14/2 module support [28-2](#)

- specifying TCP connections [28-32](#)

- virtual ISLs [28-20](#)

FCIP interfaces

- binding to port numbers [27-20](#)
- configuring advanced features [28-30 to 28-36](#)
- creating [28-30](#)

FCIP links

- B port interoperability mode [28-34](#)
- configuring [28-22](#)
- configuring peers [28-30](#)
- configuring QoS [28-36](#)
- creating [28-24](#)
- description [28-21](#)
- end points [28-21](#)
- initiating IP connections [28-32](#)
- TCP connections [28-21](#)

FCIP profiles

- configuring listener ports [28-25](#)
- configuring TCP parameters [28-25 to 28-30](#)
- creating [28-23](#)
- description [28-21](#)

FCP

- intermixing protocols [27-4](#)
- routing requests [28-53](#)

fcping

- invoking [39-7](#)

FCS

- configuring [40-2](#)
- description [40-1](#)
- logging facility [36-2](#)
- significance [40-2](#)

fctrace

- invoking [39-5](#)

FDMI

- displaying [18-5](#)

Fibre Channel

- iSCSI targets [28-54 to 28-58](#)

Fibre Channel analyzers [38-10](#)

Fibre Channel Congestion Control. See FCC

Send documentation comments to mdsfeedback-doc@cisco.com.

Fibre Channel domain. See [fcdomain](#)

Fibre Channel interface

default settings [12-33](#)

Fibre Channel interfaces

characteristics [12-2 to 12-19](#)

configuring [12-9](#)

configuring auto mode [12-10](#)

configuring BB_credits [12-12](#)

configuring beacon modes [12-16](#)

configuring data field sizes [12-16](#)

configuring extended BB_credits [12-14](#)

configuring frame encapsulation [12-16](#)

configuring performance buffers [12-13](#)

configuring port mode [12-10](#)

displaying VSAN membership [10-10](#)

modes [12-3 to 12-5, 12-10](#)

states [12-6 to 12-8](#)

Fibre Channel over IP. See [FCIP](#)

Fibre Channel PortChannels

redundancy [28-50](#)

Fibre Channel Protocol. See [FCP](#)

Fibre Channel targets

dynamic importing [28-55](#)

Fibre Channel traffic

SPAN sources [38-4](#)

Fibre Channel write acceleration

default settings [25-21](#)

description [25-5](#)

displaying configuration [25-7](#)

enabling [25-5](#)

FICON

advantages [27-3](#)

automatic save [27-19](#)

configuration files [27-23](#)

configuring [27-1](#)

default settings [27-48](#)

FC4 protocols [27-2](#)

FCIP support [27-4](#)

MDS-supported features [27-5](#)

PortChannel support [27-4](#)

port numbering [27-7](#)

prohibited ports [27-33](#)

setting up [27-12](#)

VSAN offline state [27-27](#)

files

copying [2-25, 4-28](#)

deleting [4-33](#)

displaying checksums [2-23](#)

displaying contents [2-25](#)

displaying last lines [2-27](#)

moving [2-24](#)

file systems

accessing standby supervisor modules [4-33](#)

creating directories [2-23](#)

deleting directories [2-24](#)

displaying current directory [2-23](#)

error recovery [6-30, 6-32](#)

formatting [2-21](#)

listing files [2-23](#)

redirection [2-26](#)

setting current directory [2-22](#)

volatile: [2-20](#)

File Transfer Protocol. See [FTP](#)

filters

capture [39-17](#)

defining display [39-14](#)

Flash devices

external CompactFlash [2-20](#)

formatting [2-21](#)

internal bootflash: [2-20](#)

overview [2-20](#)

FLOGI

displaying details [18-1](#)

logging facility [36-2](#)

flow statistics [24-15](#)

FL port mode

classes of service [12-4](#)

description [12-4](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

FL ports

- configuring [12-10](#)
- fctrace [39-5](#)
- nonparticipating code [12-8](#)
- persistent FC IDs [31-10](#)
- SPAN [38-4](#)

F port mode

- classes of service [12-4](#)
- description [12-4](#)

F ports

- configuring [12-10](#)
- SPAN [38-4](#)

FPSF

- load balancing [28-49](#)

frame encapsulation

- configuring [12-16](#)

frames

- configuring MTU size [28-6](#)
- encapsulation [38-8](#)
- reordering [24-11](#)

FSPF

- alternative paths [24-1](#)
- clearing counters [24-9](#)
- computing link cost [24-6](#)
- configuring globally [24-4](#)
- configuring on interfaces [24-6](#)
- default settings [24-21, 26-26](#)
- disabling on interfaces [24-7](#)
- disabling routing protocols [24-5](#)
- hello time intervals [24-6](#)
- hold time range [24-1](#)
- interoperability [39-23](#)
- link state protocol [24-2](#)
- reconvergence time [24-2](#)
- routing services [24-1](#)
- topologies example [24-2](#)

FTP

- logging facility [36-2, 36-7](#)

full core dumps

- IPS modules [28-120](#)

full zones

- default settings [15-36](#)

full zone sets

- considerations [15-8](#)
- distributing [15-12](#)

Fx ports

- 32-port default [12-8](#)
- configuring [12-10](#)
- description [12-5](#)
- FCS [40-1](#)
- interface modes [12-5](#)

G

Gigabit Ethernet

- default parameters [28-121](#)

Gigabit Ethernet interfaces

- configuring [28-4 to 28-19](#)
- configuring auto-negotiation [28-6](#)
- configuring high availability [28-15 to 28-19](#)
- configuring MTU frame size [28-6](#)
- configuring promiscuous mode [28-6](#)
- configuring static IP routing [28-8](#)
- displaying statistics [28-12 to 28-15](#)
- subinterfaces [28-8](#)
- subnet requirements [28-8](#)

Gigabit Ethernet subinterfaces

- configuring VLANs [28-7](#)

graceful shutdown

- description [1-5](#)

guidelines

- port swapping [27-26](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

H

hardware

- displaying inventory [8-2](#)
- displaying temperature [8-10](#)

hard zoning

- description [15-10](#)

HBA ports

- configuring area FCIDs [31-11](#)

HBAs

- device aliases [16-1](#)

Hello time intervals

- configuring [24-6](#)

hidden routes [24-17](#)

high availability

- description [1-4, 5-1](#)
- displaying status [5-5](#)
- Ethernet PortChannel [28-93](#)
- Ethernet PortChannels [28-50](#)
- Fibre Channel PortChannels [28-50](#)
- Gigabit Ethernet features [28-15](#)
- licensing [3-5](#)
- process restartability [5-4](#)
- protection against link failure [5-1](#)
- software upgrade [6-4](#)
- switchover characteristics [5-2](#)
- VRRP [28-49, 28-92](#)

host control

- FICON [27-17](#)

customer IDs [30-3](#)

image version and IDs [6-2](#)

login IDs [4-6](#)

process IDs [4-29, 41-1, 41-6](#)

profile IDs [30-4](#)

region ID [24-5](#)

serial IDs [30-22](#)

server IDs [30-22](#)

site IDs [30-3, 30-21](#)

See also destination IDs

See also exchange IDs

See also port IDs

See also source IDs

See also user IDs

See also VR IDs

IKE

algorithms [29-6](#)

configuring an IPsec domain [29-7](#)

default settings [29-36](#)

description [29-3](#)

initializing [29-7](#)

refreshing SAs [29-12](#)

transforms [29-6](#)

IKE domains

clearing [29-11](#)

IKE policies

configuring parameters [29-9](#)

negotiation [29-8](#)

IKE tunnels

clearing [29-11](#)

description [29-8](#)

images

See kickstart images; software images; system images

implemented port [27-8](#)

in-band management

CUP [27-27](#)

ingress port [38-11](#)

in-order delivery [24-11](#)

enabling [24-10, 24-12, 24-13](#)

ICMP packets

- type value [26-7](#)

ICMP statistics

- displaying [28-15](#)

IDs

- Cisco.com IDs [30-2](#)
- contract IDs [30-3, 30-21](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

- in-order guarantee [24-11](#)
- install all
 - command examples [6-9](#)
- install all command
 - benefits [6-6](#)
 - examples [6-9, 6-12](#)
 - failure cases [6-6](#)
 - function [6-5](#)
 - remote location path (caution) [6-14](#)
 - requirements [6-3](#)
 - usage [6-7](#)
- installed port [27-9](#)
- Intelligent Storage Services
 - configuring SSI boot variables [7-21](#)
- interfaces
 - adding to PortChannels [14-8](#)
 - administrative states [12-6](#)
 - configuring [12-9](#)
 - configuring data field size [12-16](#)
 - configuring descriptions [12-12](#)
 - configuring FSPF [24-6](#)
 - configuring zone members [15-5](#)
 - default settings [12-33](#)
 - displaying information [12-21 to 12-31](#)
 - isolated states [14-10](#)
 - modes [12-10](#)
 - operational states [12-6](#)
 - reason codes [12-6](#)
 - suspended states [14-10](#)
 - troubleshooting operational states [12-7](#)
- internal bootflash:
 - description [2-20](#)
 - initializing [2-21](#)
 - kickstart images [2-21](#)
 - recovering from corruption [2-21](#)
 - system images [2-21](#)
- internal switch states
 - description [5-6](#)
- Internet Key Exchange. See IKE
- interoperability
 - configuring [39-22](#)
 - verifying status [39-25](#)
- Inter-Switch Links. See ISL
- Inter-VSAN Routing. See IVR
- Inter-VSAN zones. See IVZs
- Inter-VSAN zone sets. See IVZSs
- invoking fcping [39-7](#)
- IOD. See in-order delivery
- IP Access Control Lists. See ACLs
- IP addresses
 - configuring in VSANs [26-12](#)
 - SMTP server [30-10](#)
- IP connections
 - active mode [28-32](#)
 - initiating [28-32](#)
- IPFC
 - logging facility [36-2](#)
- IP forwarding
 - disabling [26-12](#)
- IP over Fibre Channel. See IPFC
- IP routing
 - configuring on Gigabit Ethernet interfaces [28-8](#)
 - displaying the route table [28-9](#)
 - static [1-8](#)
- IPS core dumps. See core dumps
- IPsec
 - crypto ACLs [29-12 to 29-16](#)
- IPsec
 - algorithms [29-6](#)
 - compatibility [29-4](#)
 - configuring [29-12](#)
 - default settings [29-36](#)
 - description [29-2](#)
 - prerequisites [29-3](#)
 - transforms [29-6](#)
- IP Security. See IPsec
- IP services
 - default settings [26-26](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

IPS modules

- CDP support [28-19](#)
- configuring CDP [4-40](#)
- core dumps [28-120](#)
- partial core dumps [28-120](#)
- port modes [28-5](#)
- supported features [28-2](#)

IPS port mode

- description [28-5](#)

IPS ports

- multiple connections [28-90](#)

IP statistics

- displaying [28-13](#)

IP Storage services modules. See IPS modules

IQN

- formats [28-54](#)

iSCSI

- access control [28-67 to 28-70](#)
- configuring [28-51 to 28-93](#)
- configuring VRRP [28-92](#)
- default parameters [28-122](#)
- displaying global information [28-80](#)
- drivers [28-51](#)
- Gigabit Ethernet ports [28-5](#)
- IPS module support [28-2](#)
- MPS-14/2 module support [28-2](#)
- PortChannel-based high availability [28-93](#)
- requests and responses [28-53](#)

iSCSI authentication

- configuring [28-70](#)
- global override [28-71](#)
- scenarios [28-93](#)
- setup guidelines [28-93](#)

iSCSI high availability

- configuring [28-86 to 28-93](#)

iSCSI hosts

- VSAN membership [28-65](#)

iSCSI hosts

- initiator identification [28-59](#)
- initiator presentation modes [28-60](#)

iSCSI initiators

- assigning WWNs [28-62](#)
- configuring static IP address mapping [28-62](#)
- displaying information [28-82 to 28-85](#)

iSCSI interfaces

- configuring [28-59 to 28-76](#)
- configuring listener ports [28-74](#)
- configuring routing mode [28-75](#)
- configuring TCP tuning parameters [28-74](#)
- creating [28-53](#)
- displaying information [28-77](#)
- VSAN membership [28-66](#)

iSCSI proxy initiators

- displaying information [28-79](#)

iSCSI sessions

- authentication [28-70 to 28-73](#)
- displaying information [28-81](#)

iSCSI targets

- advertising [28-57](#)
- dynamic importing [28-54](#)
- examples [28-57](#)
- secondary access [28-88](#)
- static importing [28-56](#)
- transparent failover [28-86](#)

iSCSI users

- displaying information [28-86](#)

iSCSI virtual targets

- displaying information [28-86](#)

ISLs

- PortChannel links [14-2](#)

iSMS servers

- enabling [28-112](#)

iSNS

- configuring [28-106 to 28-113](#)
- displaying configurations [28-114 to 28-120](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

iSNS profiles

creating [28-107](#)

iSNS servers

configuration distribution [28-112](#)

isolated VSAN

description [10-5](#)

displaying membership [10-5](#)

IVR

activating topologies [17-12](#)

adding virtual domains [17-15](#)

AF IDs [17-10](#)

border switches [17-4](#)

configuration distribution with CFS [17-6](#)

configuring [17-5](#)

configuring logging levels [17-25](#)

database merge guidelines [17-23](#)

default settings [17-31](#)

description [17-2](#)

displaying information [17-25 to 17-28](#)

edge switches [17-4](#)

edge VSANs [17-3](#)

enabling [17-6](#)

features [17-3](#)

interoperability [17-22](#)

NAT [17-9](#)

native VSANs [17-3](#)

paths [17-3](#)

read-only zoning [17-22](#)

sample configuration [17-28 to 17-31](#)

sharing resources [17-2](#)

terminology [17-3](#)

transit VSANs [17-3](#)

unique domain ID guidelines [17-4](#)

zone communication [17-16](#)

zones [17-3](#)

zonesets [17-3](#)

IVR service groups

configuring [17-22](#)

description [17-21](#)

IVR topologies

clearing [17-14](#)

configuring [17-10 to 17-13](#)

configuring automatic discovery [17-13](#)

description [17-10](#)

manually creating [17-10](#)

migrating from automatic mode to user-configured mode [17-13](#)

IVZs

automatic creation [17-16](#)

clearing database [17-20](#)

configuring [17-18 to 17-21](#)

configuring QoS attributes [17-19](#)

description [17-3, 17-16](#)

differences with zones (table) [17-16](#)

displaying configurations [17-21](#)

LUN zoning [17-19](#)

IVZSs

activating with force option [17-20](#)

configuring [17-18 to 17-21](#)

description [17-3, 17-16](#)

J

jitter

configuring estimated maximum in FCIP profiles [28-29](#)

jumbo frames. See MTUs

K

keepalive timeouts

configuring in FCIP profiles [28-26](#)

kernel core dumps [41-8](#)

configuring [41-9](#)

kickstart images

KICKSTART variable [6-1](#)

loading system images [6-25](#)

overview [6-2](#)

recovering corrupted [6-30](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

recovery [6-31](#)

recovery interruption [6-26](#)

L

last byte

FC IDs [27-16](#)

LEDs

link [12-17](#)

speed [12-17](#)

status [12-17](#)

libpcap freeware [39-8](#)

licenses

Cisco MDS 9216i switches [3-3](#)

description [3-1](#)

displaying information [3-12](#)

enterprise package [3-4](#)

expiry alerts [3-11](#)

extended BB_credits [12-14](#)

Fabric Manager Server package [3-5](#)

factory-installed [3-6](#)

feature-based [3-3](#)

high availability [3-5](#)

identifying features in use [3-9 to 3-11](#)

installation options [3-6](#)

installing key files [3-8](#)

installing manually [3-6](#)

mainframe package [3-5](#)

module-based [3-3](#)

obtaining key files [3-7](#)

SAN extension package [3-4](#)

standard package [3-3](#)

Storage Services Enabler package [3-5](#)

terminology [3-2](#)

transferring between switches [3-12](#)

licensing

description [1-4](#)

link cost [24-2](#)

link failure

protection against [5-1](#)

link redundancy

Ethernet PortChannels [28-17](#)

load balancing

attributes [10-6](#)

description [14-4](#)

FSPF [28-49](#)

guarantee [10-7](#)

PortChannel [28-48](#)

PortChannels [14-2](#)

loader

loading kickstart [6-25](#)

local capture [39-10](#)

locking mechanism

FICON files [27-23](#)

log files [41-6](#)

configuring [36-5](#)

logging

default settings [36-13](#)

severity levels [36-3](#)

system messages [36-1](#)

logical unit numbers. See LUNs

loop monitoring [39-22](#)

loop port [39-22](#)

LSR [24-19, 24-20](#)

LUNs

displaying discovered, example [37-4](#)

IVR zoning [17-19](#)

LUN zoning

description [15-17](#)

M

MAC= keyword [26-9](#)

mainframe

FICON parameters [27-17](#)

VSAN clock [27-17](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

- mainframe package licenses
 - description [3-5](#)
- management
 - redundancy [5-2](#)
- management access
 - description [4-14](#)
 - in-band [4-4, 4-9 to 4-13](#)
 - obtaining [4-22](#)
 - out-of-band [4-4, 4-5 to 4-9](#)
 - using force option [4-22](#)
- Management Information Base. See MIB
- management interfaces
 - configuring [12-19](#)
 - default settings [12-33](#)
 - features [12-19](#)
- management interfaces. See mgmt0 interfaces
- manually enabling
 - FICON [27-15](#)
- maximum retransmissions
 - configuring in FCIP profiles [28-26](#)
- MD5 authentication [26-23](#)
- memory test [6-26, 6-27](#)
- mgmt0 interfaces
 - autosensing port [4-21](#)
 - configuring [4-21, 12-19](#)
 - configuring ethernet ports [26-3](#)
 - configuring speed [12-11](#)
 - default settings [12-33](#)
 - features [12-19](#)
 - recovery from switch(boot)# prompt [6-31](#)
- minimum retransmit timeouts
 - configuring in FCIP profiles [28-26](#)
- modem connections
 - configuration guidelines [4-36](#)
 - configuring [4-36 to 4-39](#)
 - configuring default initialization strings [4-38](#)
 - configuring user-specified initialization strings [4-38](#)
 - enabling [4-36](#)
 - initialization strings [4-37](#)
 - initializing on power-on switch [4-39](#)
 - verifying configuration [4-39](#)
- module configuration
 - purging [7-8](#)
 - saving to NVRAM [7-7](#)
- modules
 - configuring logging [36-4](#)
 - displaying temperatures [8-10](#)
 - preserving the configuration [7-7](#)
 - purging configurations [7-8](#)
 - replacing [6-24](#)
 - resetting [7-6](#)
 - state descriptions [7-4](#)
 - temperature monitoring [8-9](#)
 - verifying status [4-15, 7-3](#)
- monitoring traffic [38-7, 38-17](#)
- MPS-14/2 module
 - configuring extended BB_credits [12-14](#)
 - functions [6-9](#)
 - overview [1-3](#)
- MPS-14/2 modules
 - CDP support [28-19](#)
 - port modes [28-5](#)
 - supported features [28-2](#)
- MTU frame size
 - configuring Gigabit Ethernet interfaces [28-6](#)
- MTUs
 - configuring size
- multicast routing [24-10](#)
- Multiprotocol Services module. See MPS-14/2 module
- Multiprotocol Services modules. See MPS-14/2 modules
- mutual CHAP authentication
 - configuring for iSCSI [28-72](#)

N

- name server
 - interoperability [39-24](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

name servers

- displaying database [18-3](#)
- proxy feature [18-2](#)
- registering proxies [18-3](#)

NAT

- description [17-8](#)
- enabling [17-9](#)

Network Address Translation. See NAT

network administrators

- additional roles [19-3](#)
- permissions [2-3, 19-3](#)

network operators

- permissions [2-3, 19-3](#)

Network Time Protocol. See NTP

network traffic

- monitoring [38-7, 38-17](#)

next hop domain ID [24-9](#)

NL ports

- fctrace [39-5](#)
- interface modes [12-5](#)
- zone enforcement [15-10](#)

nondisruptive

- upgrades [6-4](#)

nonparticipating codes

- description [12-8](#)

N ports

- fctrace [39-5](#)
- zone enforcement [15-10](#)
- zone membership [15-2](#)

NTP

- committing configuration changes [4-20](#)
- configuration guidelines [4-18](#)
- configuring [4-17](#)
- configuring CFS distribution [4-19](#)
- database merge guidelines [4-21](#)
- discarding configuration changes [4-20](#)
- logging facility [36-2](#)
- releasing fabric session lock [4-21](#)

time-stamp option [28-33](#)

verifying session status [4-21](#)

Nx ports

hard zoning [15-10](#)

O

of [28-36](#)

one-step upgrade

- install all command [6-4](#)
- reload command [6-4](#)

operational interfaces

- viewing PortChannels [14-17](#)

operational states

- description [12-6](#)
- setting [12-10](#)

originator exchange IDs. See exchange IDs

out-of-order delivery [24-11](#)

P

packets

- discarding in FCIP [28-32](#)

passive mode

- IP connection [28-32](#)

password recovery [19-37](#)

passwords

- administrator [4-3](#)
- setting administrator default [4-5, 4-10](#)

path discovery [39-5](#)

path MTU. See PMTU

performance buffers

- configuring [12-13](#)

periodic inventory notification [30-11](#)

permit conditions [26-5, 26-8](#)

permitted filters [39-17](#)

persistent domain ID

- FICON VSANs [27-39](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

- persistent FC IDs
 - description [31-9](#)
 - displaying [31-15](#)
- PMTUs
 - configuring in FCIP profiles [28-27](#)
- port addresses
 - assigning names [27-22](#)
 - FICON [27-8](#)
- port blocking [27-21](#)
- PortChannel
 - interfaces [28-57](#)
 - subinterfaces [28-57](#)
- PortChannels
 - 32-port switching module configuration
 - guidelines [14-3](#)
 - adding interfaces [14-8](#)
 - binding to port number [27-20](#)
 - comparison with trunking [14-4](#)
 - compatibility checks [14-10](#)
 - configuration guidelines [14-11](#)
 - configuring [14-6](#)
 - configuring FC routes [24-9](#)
 - configuring for FCIP high availability [28-48](#)
 - default settings [14-20](#)
 - deleting [14-8](#)
 - description [1-8, 14-2](#)
 - examples [14-2](#)
 - FICON support [27-4](#)
 - forcing interface additions [14-9](#)
 - guidelines [14-11](#)
 - high availability [5-1](#)
 - in-order guarantee [24-12](#)
 - interoperability [39-23](#)
 - IQN formats [28-55](#)
 - link changes [24-11](#)
 - link failures [24-3](#)
 - load balancing [1-8, 14-4, 28-48](#)
 - logging facility [36-2](#)
 - member combinations [28-17](#)
 - misconfiguration error detection [14-11](#)
 - misconfiguration examples [14-13](#)
 - reason codes [12-8](#)
 - SPAN [38-4](#)
 - valid example configurations [14-12](#)
- port groups
 - 16-port switching modules [12-17](#)
 - assigning extended BB_credits [12-14](#)
- port IDs
 - configuring zone members [15-4](#)
- port modes
 - auto [12-5](#)
 - description [12-3 to 12-5](#)
 - IPS [28-5](#)
- port numbers
 - binding to FCIP interfaces [27-20](#)
 - binding to PortChannels [27-20](#)
 - FICON [27-8](#)
- ports
 - aggregation [5-1](#)
 - prohibiting [27-21](#)
 - virtual E [28-20](#)
 - VSAN membership [10-6](#)
- port security
 - default settings [21-18](#)
 - fabric binding comparison [27-37](#)
- port swapping
 - FICON [27-25](#)
 - guidelines [27-26](#)
- port tracking
 - default settings [33-8, 35-9](#)
 - description [1-14](#)
 - displaying information [33-6, 35-8](#)
 - overview [33-1](#)
- port world wide names. See pWWNs
- power supplies
 - configuring [8-6](#)
 - default state [8-12](#)
 - displaying configuration [8-6](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

- guidelines [8-6 to 8-8](#)
- modes [4-30](#)
- power usage
 - displaying [8-5](#)
- preferred domain IDs [31-5](#)
- preshared key [19-6](#)
- principal switches
 - assigning domain ID [31-4](#)
 - configuring [31-5](#)
 - selecting [31-1](#)
- private device [12-30](#)
- processes
 - nondisruptive restarts [5-1](#)
 - restartability [5-4](#)
- process ID [41-6](#)
- Process Logs [41-3](#)
- process restartability [5-4](#)
- prohibited ports
 - FICON [27-33](#)
- promiscuous mode
 - configuring Gigabit Ethernet interfaces [28-6](#)
- protocol analysis [39-8](#)
- proxies
 - registering [18-3](#)
- pWWNs
 - configuring zone members [15-4](#)
 - rejecting duplicates [18-3](#)
 - zone membership [15-2](#)

Q

- QoS
 - default settings [32-14](#)
 - displaying information [32-4, 32-10](#)
 - DSCP value [28-36](#)
 - enabling [32-6](#)
 - enabling control traffic [32-4](#)

- logging facilities [36-2](#)
- priority queuing [1-10](#)
- quality of service. See QoS

R

RADIUS

- AAA solutions [1-13, 19-1](#)
- configured parameters [19-9](#)
- secret key [1-13, 19-1](#)
- setting preshared key [19-6](#)
- specifying servers [19-6](#)
- specifying time-out [19-7](#)
- read-only zones
 - configuration guidelines [15-19](#)
 - default settings [15-36](#)
 - description [15-19](#)
- reconfigure fabric frames
 - description [31-3](#)
- reconvergence time
 - FSPF [24-2](#)
- recovering passwords [19-37](#)
- recovery sequence [6-26](#)
- redundancy
 - Ethernet PortChannels [28-50](#)
 - Fibre Channel PortChannels [28-50](#)
 - VRRP [28-49](#)
- redundancy states
 - value descriptions [5-6](#)
- redundant physical links [24-3](#)
- Registered Link Incident Report. See RLIR
- Registered State Change Notifications. See RSCNs
- reloading
 - modules [7-20, 7-22](#)
- remote capture [39-10, 39-12](#)
- remote capture daemon [39-9](#)
- Remote Capture Protocol. See RPCAP
- Remote Monitoring. See RMON
- Remote SPAN. See RSPAN

Send documentation comments to mdsfeedback-doc@cisco.com.

retransmit intervals [24-7](#)

RLIR

FICON-enabled switches [27-44](#)

Sending LIRs [27-1](#)

RMON

default settings [23-3](#)

displaying information [23-3](#)

role-based access

description [1-13](#)

roles

defaults [2-3](#)

route cost

computing [24-6](#)

routing

multicast [24-10](#)

See also broadcast routing

See also IP routing

RPCAP

Ethereal communication [39-10](#)

rsa1 key pairs

generating [19-35](#)

rsa key pairs

generating [19-35](#)

RSCN

logging facility [36-2](#)

RSCNs

displaying notifications [18-7](#)

multiple port IDs [18-8](#)

RSPAN

default settings [38-30](#)

description [1-11](#)

running configuration files

saving to startup configuration file [4-27](#)

run time checks [24-8](#)

S

SACKs

configuring in FCIP profiles [28-27](#)

SAN extension package licenses

description [3-4](#)

SAN extension tuner

assigning SCSI read/write commands [34-5](#)

configuring [34-3](#)

configuring nWWNs [34-4](#)

configuring virtual N ports [34-5](#)

data patterns [34-7](#)

default settings [34-9](#)

description [1-14, 34-2](#)

displaying tuning configuration [34-8](#)

initialization [34-3](#)

license requirements [34-2](#)

tuning guidelines [34-3](#)

SAN operating system. See Cisco MDS SAN-OS

SAN Tap

description [25-10 to 25-14](#)

displaying information [25-15, 25-20](#)

enabling [25-14](#)

proxy mode-1 [25-13](#)

proxy mode-2 [25-14](#)

transparent mode [25-12](#)

SCP

copying images [6-18](#)

scripts

setting delay time [2-28](#)

SCSI

routing requests [28-52](#)

SCSI flow configuration client

description [25-3](#)

SCSI flow data path support

description [25-3](#)

SCSI flow manager

description [25-3](#)

SCSI Flow Services

configuring [25-3](#)

default settings [25-21](#)

description [25-2](#)

displaying [25-7](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

- Fibre Channel write acceleration [25-5](#)
- functional architecture (figure) [25-2](#)
- SCSI flow configuration client [25-3](#)
- SCSI flow data path support [25-3](#)
- SCSI flow manager [25-3](#)
- SCSI flow statistics
 - default settings [25-21](#)
 - description [25-6](#)
 - displaying [25-7](#)
 - enabling [25-7](#)
- SCSI LUNs
 - discovering targets [37-1](#)
- SD port mode
 - description [12-5](#)
 - interface modes [12-5](#)
- SD ports
 - bidirectional traffic [38-12](#)
 - configuring [12-10, 38-7, 38-22](#)
- secondary MAC address [39-19](#)
- secure access
 - description [1-13](#)
- Secure Shell. See SSH
- security control
 - local [19-2, 19-19](#)
 - remote [19-2, 19-5, 19-10](#)
- security management
 - description [1-12](#)
- security parameter index. See SPI
- selective acknowledgments. See SACKs
- selective purging
 - persistent FC IDs [31-13](#)
- serial numbers
 - displaying [8-4](#)
- services module
 - purging configurations [7-8](#)
- services modules
 - description [7-3](#)
 - managing [7-1](#)
 - monitoring states [7-1](#)
 - power cycling [7-6](#)
 - replacing [6-24](#)
 - resetting [7-6](#)
 - state descriptions [7-4](#)
 - verifying status [7-3](#)
- severity levels
 - logging [36-5](#)
- SFPs
 - transmitter types [12-18](#)
- Simple Network Management Protocol. See SNMP
- simple text authentication [26-23](#)
- simulating
 - Call Home [30-14](#)
- slot0:
 - formatting [2-21](#)
- small computer system interface. See SCSI
- SMARTnet [30-2](#)
- SMTP
 - server address [30-10](#)
- SNMP
 - access control [22-2](#)
 - access groups [22-4](#)
 - adding communities [22-9](#)
 - community strings [22-2](#)
 - configuring from CLI [22-7](#)
 - counter Information [22-14](#)
 - creating roles [22-5](#)
 - creating users [22-6](#)
 - default groups [22-6](#)
 - default settings [22-15](#)
 - displaying information [22-13, 23-3](#)
 - FICON control [27-18](#)
 - read-write access [22-9](#)
 - server contact [30-3](#)
 - SNMPv3 access [1-13](#)
 - Version 3 security features [22-2](#)
 - versions supported [22-2](#)
- SNMP manager
 - FCS [40-2](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

- software images
 - bootflash corruption [6-25](#)
 - compatibility issues [6-19](#)
 - corruption [6-25](#)
 - default setting [6-34](#)
 - error state [6-25](#)
 - recognizing errors [6-34](#)
 - recovery procedure [6-25](#)
 - space requirement [6-3](#)
 - synchronizing [5-4](#)
 - upgrade requirements [6-2](#)
 - upgrading [6-1](#)
 - variables [6-1](#)
- software upgrades
 - disruptive [6-7](#)
 - manual, dual supervisor [6-18](#)
 - mechanisms [6-4](#)
 - nondisruptive [5-1](#)
 - quick [6-23](#)
- soft zoning
 - description [15-10](#)
- source IDs
 - Call Home event format [30-22](#)
 - exchange based [14-6](#)
 - flow based [14-5](#)
 - frame identification [32-2](#)
 - frame loop back [39-5](#)
 - in-order delivery [24-11](#)
 - load balancing [1-8](#)
 - path selection [10-6](#)
- SPAN
 - configuring sessions [38-5](#)
 - default settings [38-29](#)
 - description [1-11](#)
 - egress source [38-3](#)
 - encapsulating frames [38-8](#)
 - FC analyzers [38-10](#)
 - ingress source [38-3](#)
 - monitoring traffic [1-11, 38-2](#)
 - source configuration [38-4](#)
 - sources [38-4](#)
- SPAN destination port mode. See SD port mode
- SPAN tunnel port mode. See ST port mode
- special frames
 - enabling [28-31](#)
- SPI
 - configuring virtual router [26-23](#)
- SSH
 - default service [19-34](#)
 - host key pair [19-35](#)
 - protocol status [19-37](#)
 - session [6-18](#)
- SSH key pair
 - overwriting [19-35](#)
- SSH session
 - message logging [36-4](#)
- SSI boot variables
 - configuring ASMs for switching [7-17](#)
 - configuring SSMs for Intelligent Storage Services [7-21](#)
 - configuring SSMs for switching [7-21](#)
- SSI variables
 - specifying [7-21](#)
- SSMs
 - configuring Intelligent Storage Services [25-1 to 25-21](#)
 - overview [1-3](#)
 - reloading [7-22](#)
 - specifying boot images [7-21](#)
 - specifying image boot variables [7-18 to 7-22](#)
- standard package licenses
 - description [3-3](#)
- standby modules
 - monitoring [5-2](#)
- standby supervisor modules
 - accessing file systems [4-33](#)
 - copying boot variables [5-4](#)
 - synchronizing [5-4](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

- startup configuration files
 - saving running configuration file [4-27](#)
 - unlocking [4-28](#)
- static domain IDs [31-5](#)
- static routes
 - run time checks [24-8](#)
- storage devices
 - access control [15-1](#)
 - permanent [2-20](#)
 - temporary [2-20](#)
- Storage Services Enabler package licenses
 - description [3-5](#)
- Storage Services Modules. See SSMs
- ST port mode
 - description [12-5](#)
- ST ports
 - configuring [38-19](#)
 - interface modes [12-5](#)
- subnet mask
 - BIOS setup configuration [6-28](#)
 - configuring IP routes [26-13](#)
 - configuring mgmt0 [4-21, 26-2](#)
 - configuring switch [4-3](#)
 - default setting [7-23](#)
 - initial configuration [4-7, 4-11](#)
 - loader> prompt recovery [6-30](#)
 - switch(boot)# prompt recovery [6-31](#)
- subnet masks
 - configuring mgmt0 interfaces [12-19](#)
- subnets
 - requirements [28-8](#)
- subordinate switch [31-7](#)
- supervisor module
 - CDP support [28-19](#)
- supervisor modules
 - active [5-2](#)
 - active mode [1-11](#)
 - active state [5-6, 7-4](#)
 - default settings [7-23](#)
 - description [7-1](#)
 - displaying information [7-5](#)
 - high availability [5-2](#)
 - manual switchovers [5-2](#)
 - recovering password [6-32](#)
 - redundancy [5-1](#)
 - resetting [7-6](#)
 - standby mode [1-11](#)
 - standby state [5-6, 7-4](#)
 - state descriptions [5-6, 7-4](#)
 - switch options [1-11](#)
 - switchover mechanisms [5-2](#)
 - switchovers after failure [5-2](#)
 - synchronizing [5-4](#)
 - verifying status [7-3](#)
- supervisors
 - replacing [6-24](#)
- Switched Port Analyzer. See SPAN
- switches
 - displaying serial numbers [8-4](#)
 - display power usage [8-5](#)
 - dual supervisor [6-32](#)
 - rebooting [7-6](#)
 - reliability service [1-5](#)
 - reloading [7-6](#)
 - single supervisor [6-31](#)
- switching modules
 - accessing [7-5](#)
 - description [7-3](#)
 - managing [7-1](#)
 - monitoring states [7-1](#)
 - power cycling [7-6](#)
 - powering off [7-8](#)
 - preserving configuration [7-7](#)
 - purging configurations [7-8](#)
 - reloading [7-6](#)
 - resetting [7-6](#)
 - state descriptions [7-4](#)
 - verifying status [7-3](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

- switch modules
 - replacing [6-24](#)
- switch names
 - assigning [4-14](#)
- switchover mechanism
 - warm state [7-4](#)
- switchovers
 - characteristics [5-2](#)
 - guidelines [5-3](#)
 - manually initiating [5-2](#)
 - supervisor modules [5-2](#)
 - VRRP [28-49](#)
- switch ports
 - configuring attribute default values [12-17](#)
- switch priority
 - configuring [31-6](#)
- switch security
 - default settings [19-42](#)
- syslog messages. See system messages
- syslogs
 - viewing [1-12](#)
- system health
 - status [41-15](#)
- system images [6-2](#)
 - reading configuration [6-25](#)
 - recovery interruption [6-26](#)
 - SYSTEM variable [6-1](#)
- system message logging server [36-1](#)
 - configuring [36-6](#)
- system messages
 - accessing [36-1](#)
 - configuring [36-3](#)
 - default settings [36-13](#)
 - displaying configuration [36-9](#)
- system processes
 - displaying [41-1](#)
 - status [41-4](#)
- system statistics
 - CPU and memory [41-5](#)

T

- TACACS+
 - AAA solutions [1-13](#)
 - starting a distribution session [19-17](#)
- target disks [37-3](#)
- TCP connections
 - FCIP profiles [28-21](#)
 - specifying [28-32](#)
- TCP parameters
 - configuring in FCIP profiles [28-25 to 28-30](#)
- TCP ports
 - ACLs [26-6](#)
- TCP statistics
 - displaying [28-14](#)
- Telnet
 - default service [19-34](#)
 - session [6-18](#)
- Telnet server connections
 - disabling [4-24](#)
- telnet server connections
 - description [4-24](#)
- Telnet session
 - message logging [36-4](#)
- temperatures
 - displaying [8-10](#)
 - major thresholds [8-9](#)
 - minor thresholds [8-9](#)
 - monitoring hardware [8-9](#)
- TE port mode
 - classes of service [12-4](#)
 - description [12-4](#)
- TE ports
 - fctrace [39-6](#)
 - FSPF topology [24-2](#)
 - interoperability [39-23](#)
 - recovering from isolation [15-13](#)
 - SPAN [38-4](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

- trunking [1-8](#)
- trunking restrictions [13-1](#)
- TFTP
 - boot [6-28](#)
 - servers [6-28, 41-6](#)
- throttling
 - duplicate messages [30-12](#)
- time
 - configuring [4-16](#)
- time out value. See TOV
- time zones
 - configuring [4-16](#)
- TL port mode
 - classes of service [12-4](#)
 - description [12-4](#)
- TL ports
 - configuring [12-10](#)
 - displaying [12-30](#)
 - FCS [40-1, 40-2](#)
 - logging facility [36-2](#)
 - SPAN [38-4](#)
 - translation guidelines [12-31](#)
- TOV
 - interoperability [39-23](#)
 - ranges [39-2](#)
- translative loop port mode. See TL port mode
- trivial authentication [19-19](#)
- troubleshooting
 - collecting output [39-29](#)
 - error messages [36-1](#)
- trunk-allowed VSAN lists
 - configuring [13-3](#)
- trunking
 - comparison with PortChannels [14-4](#)
 - configuration guidelines [13-6](#)
 - default settings [13-8](#)
 - description [1-8, 13-1](#)
 - interoperability [39-23](#)

- link state [13-2](#)
- restrictions [13-1](#)
- trunking E port mode. See TE port mode
- trunking ports
 - associated with VSANs [10-6](#)
- trunking protocol
 - default [13-2](#)
 - default settings [13-8](#)
 - description [13-2](#)
- trunk mode
 - administrative default [12-18](#)
 - default settings [13-8](#)
 - status [13-2](#)
- trunk ports
 - displaying information [13-7](#)

U

- UDP ports
 - ACLs [26-6](#)
- unblocking ports [27-27](#)
- unimplemented port [27-8](#)
- uninstalled ports [27-9](#)
- uninstalling
 - permanent licenses [3-9](#)
- updating
 - licenses [3-10](#)
- upgrades
 - disruptive [6-23](#)
 - See also disruptive upgrades
 - See also nondisruptive upgrades
- upgrading
 - software [6-18 to 6-23](#)
- upgrading BIOS. See BIOS upgrades
- user accounts
 - creating additional [4-6](#)
- user authentication
 - description [1-13](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

user IDs

- authentication [19-2](#)
- security management [1-13](#)

user profiles [19-3](#)

users

- SNMP support [22-6](#)

V

VE ports

- description [28-20](#)

virtual devices [12-30](#)

virtual E ports. See VE ports [28-20](#)

virtual ISLs

- description [28-20](#)

Virtual LANs. See VLANs

Virtual Router Redundancy Protocol. See VRRP

virtual SANs. See VSANs

VLANs

- configuring on Gigabit Ethernet subinterfaces [28-7](#)
- description [28-7](#)

volatile:

- description [2-20](#)
- switch reboots [2-22](#)

VR IDs

- configuring [26-20](#)
- mapping [26-19](#)

VRRP

- characteristics [26-19](#)
- clearing statistics [26-25](#)
- configuring for Gigabit Ethernet interfaces [28-16](#)
- configuring Gigabit Ethernet [28-16](#)
- description [28-16](#)
- group members [28-16](#)
- IQN formats [28-55](#)
- logging facility [36-3](#)
- master and backup [26-19](#)
- primary IP [26-21](#)
- priority tracking [26-23](#)

security authentication [26-23](#)

- setting priority [26-21](#)
- tracking priority [26-23](#)

VSA

- communicating attributes [19-7](#)
- protocol options [19-8](#)

VSAN IDs

- allowed list [13-8](#)
- attributes [10-6](#)
- configuring FICON [27-4](#)
- membership [10-4](#)
- multiplexing traffic [12-4](#)
- range [10-5](#)

VSAN interfaces

- configuring [12-20](#)
- displaying information [12-25](#)

VSAN membership

- iSCSI hosts [28-65](#)
- iSCSI interfaces [28-66](#)

VSANs

- advantages [10-1](#)
- allowed-active [13-1](#)
- allowed list [38-4](#)
- allowed-list [13-8](#)
- broadcast address [24-10](#)
- cache contents [31-16](#)
- clock [27-17](#)
- comparison with zones (table) [10-4](#)
- configuration overview [1-6](#)
- configuring [10-7 to 10-10](#)
- configuring domains [31-1](#)
- configuring FSPF [24-4](#)
- configuring overlay [26-17](#)
- configuring trunk-allowed lists [13-3](#)
- default settings [10-10](#)
- default VSAN [10-5](#)
- deleting [10-8](#)
- description [1-6, 10-2 to 10-6](#)
- displaying membership [10-10](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

displaying usage [10-9](#)
 domain IDs [31-7](#)
 edge [17-3](#)
 fabric optimization [27-3](#)
 FCIDs [10-2](#)
 FCS [40-1](#)
 features [10-2](#)
 flow statistics [24-15](#)
 FSPF connectivity [24-2](#)
 gateway switch [26-4](#)
 guidelines for transit VSANs [17-5](#)
 interface [12-33](#)
 interop mode [39-23](#)
 IP addresses [26-12](#)
 IPFC interface [39-6](#)
 isolated VSAN [10-5](#)
 loop devices [12-30](#)
 merging traffic [13-6](#)
 mismatch [13-2](#)
 mismatches [12-7](#)
 multiple zones [15-8](#)
 name [10-6](#)
 name server [18-2](#)
 native [17-3](#)
 operational state [10-6](#)
 overlaid routes [26-4, 26-15](#)
 port isolation [13-6](#)
 port membership [10-6](#)
 redundancy [1-6, 10-1](#)
 Rules and features [19-23](#)
 scalability [1-6, 10-1](#)
 SPAN source [38-3, 38-4](#)
 static routing [26-13](#)
 TE port mode [12-4](#)
 TOVs [39-2](#)
 traffic isolation [1-6, 10-1, 10-3](#)
 traffic routing [26-1](#)
 transit [17-3, 17-12](#)
 trunk allowed [12-33](#)

trunk-allowed [13-1, 13-2](#)
 trunking ports [10-6](#)
 VRRP [26-19](#)
 VSAN trunking. See trunking

W

window management
 configuring in FCIP profiles [28-27](#)
 world wide names. See WWNs
 WWNs
 configuring [39-18](#)
 displaying configurations [39-19](#)
 suspended connections [12-8](#)

Z

zone aliases
 conversion to device aliases [16-5](#)
 zone attribute groups
 cloning [15-21](#)
 renaming [15-21](#)
 zones
 access control [15-7](#)
 changing from enhanced zones [15-29](#)
 cloning [15-21](#)
 comparison with device aliases (table) [16-2](#)
 comparison with VSANs (table) [10-4](#)
 configuring [15-4](#)
 configuring aliases [15-6](#)
 configuring broadcasting [15-17](#)
 default policies [15-11](#)
 default policy [15-2](#)
 default settings [15-36](#)
 description [1-7](#)
 differences with IVZs (table) [17-16](#)
 displaying information [15-21 to 15-27](#)
 enforcing restrictions [15-10](#)

Send documentation comments to mdsfeedback-doc@cisco.com.

- example [15-3](#)
- exporting databases [15-13](#)
- features [15-2](#)
- implementation [15-4](#)
- importing databases [15-13](#)
- IVR communication [17-16](#)
- logging facility [36-3](#)
- LUN-based [15-17](#)
- merge failures [12-7](#)
- placing CUP in [27-28](#)
- read-only for IVR [17-22](#)
- renaming [15-21](#)
- See also default zones
- see also enhanced zones
- See also hard zoning
- see also LUN zoning
- see also read-only zones
- See also soft zoning
- zone servers
 - clearing database [15-15](#)
- zone sets
 - accesses between devices [1-7](#)
 - cloning [15-21](#)
 - considerations [15-8](#)
 - copying [15-14](#)
 - creating [15-7](#)
 - default settings [15-36](#)
 - displaying information [15-21 to 15-27](#)
 - distributing [15-11, 15-12](#)
 - exporting [15-14](#)
 - exporting databases [15-13](#)
 - importing [15-14](#)
 - importing databases [15-13](#)
 - one-time distribution [15-12](#)
 - recovering from isolation [15-13](#)
 - renaming [15-21](#)
 - See also active zone sets
 - See also full zone sets
- zone traffic priorities
- configuring [15-16](#)
- description [15-15](#)
- zoning based access control
 - configuring for iSCSI [28-67](#)

Send documentation comments to mdsfeedback-doc@cisco.com.